On the expressivity of the modal mu-calculus

J. C. Bradfield

Laboratory for Foundations of Computer Science
University of Edinburgh, The King's Buildings
EDINBURGH, United Kingdom, EH9 3JZ

email: jcb@dcs.ed.ac.uk

Abstract: We analyse the complexity of the sets of states, in certain classes of infinite systems, that satisfy formulae of the modal mu-calculus. Improving on some of our earlier results, we establish a strong upper bound (namely Δ^1_2). We also establish various lower bounds and restricted upper bounds, incidentally providing another proof that the mu-calculus alternation hierarchy does not collapse at level 2.

Keywords: descriptive complexity, logic in computer science, verification, temporal logic

1 Introduction.

This paper studies the complexity of the state sets, in certain classes of infinite systems, that satisfy formulae of the modal mu-calculus. This is of natural theoretical interest, and there is also a certain practical motivation.

The modal mu-calculus, or Hennessy–Milner logic with fix-points, is a popular logic for expressing temporal properties of systems. It is a simple logic with a deep theory, and accordingly has been much studied; it is also used in tools such as the Edinburgh Concurrency Workbench, and has been used in real verifications. Most implementations deal only with finite-state systems, and much of the applied work has concentrated on fast algorithms for model-checking finite systems. The author's interest, however, is in the verification of infinite systems, and with Stirling he developed a tableau based method for model-checking infinite systems [BrS92], and in subsequent work this system was implemented for systems modelled as Petri nets [Bra92]. Methods and tools for infinite systems must inevitably deal with infinite sets of states, and accordingly one must decide, in practice at least, on a notation for representing such sets. In the existing implementation, a simple notation is used, which is very far from being complete, in the sense of being adequate to express all tableaux for all properties of all systems. The question then arises, what notation *is* adequate, and one way of addressing this is to study the descriptive complexity of the sets of states denoted by mu-calculus formulae. One might hope, for example, that in the case of Petri nets first-order predicates on markings would be adequate, although such a hope would be vain. A few years ago, the author considered this matter in his thesis, but was only able to obtain trivial upper bounds (namely, second order arithmetic), and a reasonable lower bound.

One of the attributes of a mu-calculus formula is the *alternation depth*, which is, roughly speaking, the semantic nesting depth of alternating maximal and minimal fix-point operators. Formulae of alternation depth two (i.e. with one alternation) are essential to express properties such as 'infinitely often', but it is remarkably hard to find any practical example of a property requiring higher alternation depth (and it is also (unsurprisingly) hard to understand a formula of higher alternation depth). In fact, perhaps the main open problem concerning the modal mu-calculus is whether the alternation hierarchy collapses—most workers in the field believe that it does not, but the best result so far is that level 3 is not equal to level 2, proven by Arnold and Niwiński [ArN90] using a result of Rabin [Rab70]. This gives an additional motivation for this paper: descriptive set theorists have long studied fix-points, and one might hope to establish the non-collapse of the alternation hierarchy by transfer of results from descriptive set theory.

In this paper we first establish preliminary definitions, and review the results from the author's thesis. We then establish a much improved upper bound, namely that any mu-calculus formula in any 'reasonable' system has a Δ^1_2 denotation. Next we establish some further results, including a lower bound result (that there are simple alternation depth 1 formulae in 'simple' systems that have a Π^1_1 denotation), and a simple alternative proof of the existing non-collapse result. We conclude by discussing the notions of 'reasonable' and 'simple'.

2 Preliminaries.

2.1 Modal mu-calculus.

The modal mu-calculus, in positive form, has formulae Φ built up from variables Z by the booleans tt, ff, $\vee$ and $\wedge$, the modal connectives $[K]\Phi$ and $\langle K\rangle\Phi$, for a set K of *labels*, and the minimal and maximal fix-point connectives $\mu Z.\Phi(Z)$ and $\nu Z.\Phi(Z)$. We also write $[-K]\Phi$ for $[\mathscr{L} - K]\Phi$, and $[-]\Phi$ for $[\mathscr{L}]\Phi$. We use σ to range over μ and ν.

Note that we adopt the convention that the scope of binding operators (μ, ν, $\forall$, $\exists$) extends as far as possible.

Observe that negation is not in the language, but any closed mu-formula can be negated by using the usual De Morgan dualities—μ and ν are dual by $\neg\mu Z.\Phi(Z) = \nu Z.\neg\Phi(\neg Z)$. Where necessary, we shall assume that free variables can be negated just by adjusting the valuation. We shall use $\Rightarrow$ etc. freely, though we must ensure that bound variables only occur positively.

Given a labelled transition system $\mathscr{T} = (\mathscr{S}, \mathscr{L}, \longrightarrow)$, where $\mathscr{S}$ is a set of states, $\mathscr{L}$ a set of labels, and $\longrightarrow \subseteq \mathscr{S} \times \mathscr{L} \times \mathscr{S}$ is the transition relation (we write $s \xrightarrow{l} s'$), and given also a valuation $\mathscr{V}$ assigning subsets of $\mathscr{S}$ to variables, the denotation $\|\Phi\|_{\mathscr{V}} \subseteq \mathscr{S}$ of a mu-calculus formula Φ is defined in the obvious way for the variables and booleans, by $\|[K]\Phi\|_{\mathscr{V}} = \{\, s \mid \forall s'.\, \forall l \in K \,.\, s \xrightarrow{l} s' \;\Rightarrow\; s' \in \|\Phi\|_{\mathscr{V}} \,\}$, dually for $\langle K\rangle\Phi$, and for the fix-points by

$$\|\mu Z.\Phi\|_{\mathscr{V}} = \bigcap\{\, S \subseteq \mathscr{S} \mid \|\Phi\|_{\mathscr{V}[Z:=S]} \subseteq S \,\}$$

$$\|\nu Z.\Phi\|_{\mathscr{V}} = \bigcup\{\, S \subseteq \mathscr{S} \mid S \subseteq \|\Phi\|_{\mathscr{V}[Z:=S]} \,\}$$

It is often useful to think of μZ and νZ as meaning respectively finite and infinite looping from Z back to μZ (νZ) as one 'follows a path of the system through the formula'. Examples of properties expressible by the mu-calculus are 'always P', as $\nu Z.P \wedge [-]Z$, 'eventually P', as $\mu Z.P \vee \langle -\rangle Z$, and 'there is a path along which c happens infinitely often', as $\nu Y.\mu Z.\langle c\rangle Y \vee \langle -c\rangle Z$. (For the latter, we can loop around Y for ever, but each internal loop round Z must terminate.)

There is more than one definition of alternation depth in use, the most common being Emerson–Lei's [EmL86]; we shall use here the stronger notion of Niwiński [Niw86], also used by Kaivola in [Kai95]. A formula Φ is said to be in the classes Σ^μ_0 and Π^μ_0 iff it contains no fix-point operators. To form the class Σ^μ_{i+1}, take $\Sigma^\mu_i \cup \Pi^\mu_i$, and close under (i) boolean and modal combinators, (ii) $\mu Z.\Phi$, for $\Phi \in \Sigma^\mu_{i+1}$, and (iii) substitution of $\Phi \in \Sigma^\mu_{i+1}$ for a free variable of $\Psi \in \Sigma^\mu_{i+1}$ provided that no free variable of Φ is captured by Ψ; and dually for Π^μ_{i+1}. The *alternation depth* of Φ is the greatest i such that $\Phi \notin \Sigma^\mu_{i-1} \cup \Pi^\mu_{i-1}$. (The examples above are in Π^μ_1, Σ^μ_1 and Π^μ_2 respectively.)

A formula Φ is *unnested* if every fix-point subformula contains no occurrence of an outer fix-point variable; i.e. there is no semantic nesting of fix-points, even of the same sign. Classes of unnested formulae are defined by simple syntactic nesting: $\Sigma^{un}_0 = \Pi^{un}_0 =$ fix-point free formulae, and Σ^{un}_{i+1} is formed from $\Sigma^{un}_i \cup \Pi^{un}_i$ by closing under (i) booleans and modals, and (ii) $\mu Z.\Phi$ where $\Phi \in \Sigma^{un}_{i+1}$ and Z is not free in any fix-point subformula of Φ.

Given a formula $\mu Z.\Phi(Z)$ and a model, the *approximants* $\mu^\zeta Z.\Phi$ (for ζ an ordinal) are defined recursively by $\mu^\zeta Z.\Phi = \|\Phi\|_{Z:=S}$ where $S = \bigcup_{\eta<\zeta} \mu^\eta Z.\Phi$, and dually for $\nu Z.\Phi$.† Since mu-calculus formulae are monotonic in their variables, for successor ordinals we have $\mu^{\zeta+1} Z.\Phi = \|\Phi\|_{Z:=\mu^\zeta Z.\Phi}$. Further, by the standard Tarski fix-point theorem we have $\|\mu Z.\Phi\| = \bigcup_\zeta \mu^\zeta Z.\Phi$; the smallest κ such that $\|\mu Z.\Phi\| = \bigcup_{\zeta<\kappa} \mu^\zeta Z.\Phi$ is called the *closure ordinal* of $\mu Z.\Phi$.

2.2 Descriptive set theory.

Terminology and notation for material from descriptive set theory is standard (see Moschovakis' textbook [Mos80]), and is given here to establish notation.

ω denotes the (non-negative) integers, ranged over by $a, b, \ldots$. $\mathscr{N} = {}^\omega\omega$ is the set of infinite sequences of integers, ranged over by $\alpha, \beta, \ldots$; members of $\mathscr{N}$ are called *irrationals*. $\alpha(i)$ is the i'th element of α. We assume a coding (by, e.g., a diagonal mapping) of infinite sequences $\alpha_0, \alpha_1, \ldots$ of irrationals as a single irrational α, and the usual codings of finite sequences of (integers/irrationals) as a single (integer/irrational); $(\alpha)_i$ denotes the i'th element of the sequence coded by α, and the codes of sequences are denoted $\langle a_0, \ldots\rangle$. We also code sets of integers as (binary) irrationals in the obvious

† The usual definition in mu-calculus papers, following Kozen's choice in his seminal [Koz83], is slightly different, with different cases for zero, successor and limit ordinals; the effect of this is that our μ^ζ is Kozen's $\mu^{\zeta+1}$. The present definition is neater, and follows Moschovakis, whose presentation of descriptive set theory we use.

way—we sometimes use X, Y, Z for such irrationals. An irrational ε is called *recursive* if the map $n \mapsto \varepsilon(n)$ is a recursive function (in the usual sense).

$\mathscr{N}$ carries a topology which has as basic neighbourhoods the sets Y such that $\exists\, a_0, \ldots, a_n \,.\, \forall\, \alpha \in Y \,.\, \forall\, i \leq n \,.\, \alpha(i) = a_i$. This collection is easily recursively enunerable, so let $N(\mathscr{N}, 0), N(\mathscr{N}, 1), \ldots$ be some standard recursive enumeration of the basic neighbourhoods.

Given a space $\mathscr{X}$ that is a finite product of copies of ω and $\mathscr{N}$ with the product topology and induced basic neighbourhoods, a subset G of $\mathscr{X}$ is *semi-recursive* if it is of the form

$$G = \bigcup_n N(\mathscr{X}, \varepsilon(n))$$

for some recursive irrational ε. For $\mathscr{X} = \omega$, this coincides with the usual definition; for $\mathscr{X} = \mathscr{N}$, this means that to check that α *is* in G, it suffices to check that α matches one of some r.e. set of finite prefixes. This definition (of course) coincides with the computational interpretation of semi-recursive sets of irrationals (or reals): that there is a Turing machine, equipped with an oracle that provides the successive elements of α on demand, which accepts if $\alpha \in G$. A set P is *recursive* if both P and $\neg P$ are semi-recursive.

We can then define the usual Kleene arithmetical and analytical hierarchies: $\Sigma^0_0 = \Pi^0_0 = \{\text{recursive sets}\}$; P is Σ^0_{i+1} if there is an R in Π^0_i such that $P(x) \Leftrightarrow \exists\, s \,.\, R(s, x)$, and $\Pi^0_i = \neg\Sigma^0_i$; $\Sigma^1_0 = \Sigma^0_1$ and $\Pi^1_0 = \Pi^0_1$; P is Σ^1_{i+1} if there is an R in Π^1_i such that $P(x) \Leftrightarrow \exists\, \alpha \,.\, R(\alpha, x)$, and $\Pi^1_i = \neg\Sigma^1_i$. The class Δ^i_j is $\Sigma^i_j \cap \Pi^i_j$.

2.3 Petri nets etc.

The class of system with which we shall mostly be concerned is labelled Petri nets, by which we mean standard finite (i.e. with finite sets of places and transitions) place–transition nets, with transitions carrying labels from $\mathscr{L}$. A place–transition net is a set S of *places*, a set T of *transitions*, and functions ${}^\bullet(), ()^\bullet \colon T \to \wp_{\mathrm{m}}(S)$ assigning multisets of input and output places to transitions. A *marking* M is a multiset of places, representing the existence of a number of *tokens* on each place, and a transition t can fire at a marking M to give a marking M', written $M[t\rangle M'$, if ${}^\bullet t \leq M$ and $M' = M - {}^\bullet t + t^\bullet$—the so-called token game semantics.

We use the notion of (Minsky) register machines rather than Turing machines, for convenience in constructions. We notate a register machine as

$$(\{q_0, \ldots, q_n\}, \{R_1, \ldots, R_m\}, \{\delta_0, \ldots, \delta_{n-1}\})$$

where R_i are the registers, q_i are the states with q_0 being the initial state and q_n the unique halting state, and δ_i is the transition rule for state q_i: δ_i is either (i) '$R_j := R_j{+}1$; goto q_k' for some j, k, or (ii) 'if $R_j = 0$ then goto q_k else ($R_j := R_j - 1$; goto $q_{k'}$)' for some j, k, k'.

A *non-deterministic register machine* has a finite set of transition rules for each state, instead of a single rule.

3 Reprise of previous results.

The results in this section were obtained by the author in his thesis, published as the monograph [Bra91], and provide the basis for the rest of our results. (We have taken the opportunity to improve the proof of Theorem 3.)

The model-checking problem for Petri nets is undecidable, even for formulae with only one μ operator. This can be shown by a simple technique which in the rest of this paper allows us to consider non-deterministic Turing machines and Petri nets interchangeably, even though Petri nets lack Turing power.

Theorem 1. Given a register machine $\mathscr{M}$, there is a Petri net N and a formula Φ such that $N \vDash \Phi$ iff $\mathscr{M}$ halts.

Proof. If $\mathscr{M}$ is as above, we define N thus: the places are $q_0, \ldots, q_n, R_1, \ldots, R_m$. The transitions and flow relation are determined by the δ_i: if δ_i is of form (i), then there is a transition δ_i^+ with ${}^\bullet\delta_i^+ = \{q_i\}$ and $\delta_i^{+\bullet} = \{q_k, R_j\}$; and if δ_i is of form (ii), then there are transitions δ_i^0 and δ_i^- such that ${}^\bullet\delta_i^0 = \{q_i\}$, ${}^\bullet\delta_i^- = \{q_i, R_j\}$, $\delta_i^{0\bullet} = \{q_k\}$ and $\delta_i^{-\bullet} = \{q_{k'}\}$. So the branch in the decrement instruction becomes a non-deterministic choice.

We now design a formula which expresses the property that N deadlocks on all execution sequences corresponding to valid executions of $\mathscr{M}$. Let Δ^+ (resp. Δ^0, Δ^-) be the set of δ^+ (resp. δ^0, δ^-) transitions, let $\Delta^\pm = \Delta^+ \cup \Delta^-$, and $\Delta = \Delta^\pm \cup \Delta^0$, and define the mu-formula halt(N) to be

$$\mu Z.\Big(\sum_{i=0}^{n} q_i = 1\Big) \wedge [\Delta^\pm]Z \wedge ([\Delta^\pm]\mathrm{ff} \Rightarrow [\Delta^0]Z).$$

halt(N) says the net deadlocks, provided that we only look at paths on which the δ^0 transitions are 'honestly' fired, i.e. it's not possible to fire a δ^- transition instead (and so the relevant register is indeed zero). This is the required property. A rigorous proof is in [Bra91]. $\square$

Now by giving the register machine a designated input register with initial value x, and considering $\Phi_I \wedge \mathrm{halt}(N)$ where Φ_I is a (fix-point free) formula saying that $\mathscr{M}$ is in its initial state, we get a formula whose denotation is those initial markings such $\mathscr{M}$ halts on input x.

Note, incidentally, that we have made use of terms such as $(\sum_{i=0}^{n} q_i = 1)$ in Φ. According to taste, one can either view these as special variables, and assume that valuations give them the intended meaning (which we shall do), or remove them by converting them into (rather large) expressions using boxes and diamonds.

Still considering Petri nets, we can construct nets and unnested formulae whose denotations are at arbitrary levels in the arithmetical hierarchy.

Theorem 2. If S is a (unary, wlog) relation defined by an arithmetical formula ϕ, then there is a net N_S and a mu-formula Φ_S such that $\|\Phi_S\|$ is the set of markings such that all places are empty save that $q_0 = 1$ and $R_1 = n$ for some n such that $n \in S$, and furthermore, if S is Σ_l^0 (resp. Π_l^0) then Φ_S is Σ_l^{un} (resp. Π_l^{un}).

Proof. We sketch the proof; details are again in [Bra91]. Consider the case where S is Π_1^0, say $S(x) \Leftrightarrow \forall y \,.\, T(x,y)$ with T recursive. We build a net N representing the register machine that computes T; say the input registers are $R_x = R_1$ and R_y. Now we add to N a token generator which puts an arbitrary number of tokens in R_y and then starts N. To be precise, we add a place A_y, and two transitions: a_y such that ${}^\bullet a_y = \{A_y\}, a_y{}^\bullet = \{A_y, R_y\}$, and b_y such that ${}^\bullet b_y = \{A_y\}, b_y{}^\bullet = q_0$. The formula we want is then $\nu Y.[a_y]Y \wedge [b_y]\mathrm{halt}(N)$, plus an appropriate initialization term. We can repeat this operation; for $S = \exists y \,.\, T$, we prepend another token generator, and take the formula $\mu Y.\langle a_y\rangle Y \vee \Phi_T$ plus initialization. □

These lower bounds for the unnested case are in fact also upper bounds, and since part of the coding used in the proof is also used later, we will give details.

Theorem 3. For any Petri net N, and any formula Φ in Σ_l^{un} (resp. Π_l^{un}), $\|\Phi\|$ is in Σ_l^0 (resp. Π_l^0).

Proof. We shall construct by induction on the structure of Φ a recursive relation $F(m, n_1, \ldots, n_k)$, where m is a parameter representing a marking M of N, and $\sigma_i Z_i.\Phi_i$ for $1 \le i \le k$ are the fix-point subformulae of Φ, in order of decreasing height (and so that there is a minimal number of μ/ν alternations). The meaning of this relation is that $M \vDash \Phi[\sigma_i^{n_i} Z_i/\sigma_i Z_i]$; that is, we are replacing each fix-point by some finite approximant. Let $\vec{n}$ denote a sequence of parameters. We define F by induction on the number of fix-point subformulae. Given $\Phi = \sigma Z.\Psi$, where Ψ has $k-1$ fix-point subformulae, we first define a quantifier-free formula $\psi(m,n,\vec{n})$ by induction on the structure of Ψ:

(i) if Ψ is a variable Z, then $\psi(m,n,\vec{n})$ is $Z(m,n,\vec{n})$ where Z is a (second-order) variable of arithmetic;

(ii) if $\Psi = \Psi_1 \vee \Psi_2$, then $\psi(m,n,\vec{n})$ is $\psi_1(m,n,\vec{n}) \vee \psi_2(m,n,\vec{n})$, and similarly for $\wedge$;

(iii) if $\Psi = [K]\Psi_1$, then $\psi(m,n,\vec{n})$ is $\bigwedge_{t \in T_K}((\hat{t}(M) \Rightarrow \psi_1(\tilde{t}(M), n, \vec{n})$, where T_K is the set of transitions carrying labels in K, $\hat{t}$ is the (recursive) predicate interpreting 't is enabled at M', and $\tilde{t}$ is the (recursive) function such that $\tilde{t}(m)$ is the unique m' such that $M[t\rangle M'$;

(iv) dually for $\Psi = \langle K\rangle\Psi_1$;

(v) if $\Psi = \sigma_i Z_i.\Phi_i$, then $\psi(m,n,\vec{n})$ is $F_i(m,\vec{n}')$, where F_i is already defined by induction, and $\vec{n}'$ is the subsequence of $\vec{n}$ corresponding to the fix-point subformulae of Ψ.

Now we define $F(m,n,\vec{n})$ by the simple recursion

$$F(m,0,\vec{n}) = 0 \text{ if } \sigma = \mu,\ 1 \text{ if } \sigma = \nu$$

$$F(m,n+1,\vec{n}) = \psi(m,n,\vec{n})[F/Z]$$

(where $[-/-]$ denotes syntactic substitution);

Now, by a simple lemma from [Lar90], in a finite-branching transition system (such as a Petri net), the closure ordinal of any unnested formula is at most ω. Furthermore, if $\sigma^n Z.\Psi$ is a finite approximant, to determine if $M \vDash \sigma^n Z.\Psi$, we need to check the fix-point subformulae of Ψ on only finitely many markings $\{M_1, \ldots, M_p\}$ (where p is bounded by $n \times$ (modal depth of Ψ) $\times$ (branching degree of the system)). Hence if $\mu Z'.\Psi'$ is a subformula of Ψ, there exists n' (depending on n) such that $M \vDash \sigma^n Z.\Psi[\mu^{n'} Z'/\mu Z']$ iff $M \vDash \sigma^n Z.\Psi$). And if $\nu Z''.\Psi''$ is another subformula of Ψ, then $M \vDash \sigma^n Z.\Psi$ iff $M \vDash \sigma^n Z.\Psi[\nu^{n''} Z''/\nu Z'']$ for all n''.

Therefore we obtain

$$M \vDash \Phi \qquad \text{iff} \qquad Q_1 n_1 \dots Q_k n_k \,.\, F(m, n_1, \dots, n_k)$$

where Q_i is $\forall$ or $\exists$ according as σ_i is ν or μ, and F is the relation defined for Φ, which gives the result. □

For the case of arbitrary formulae, we have the following trivial upper bound:

Proposition 4. For any Petri net (or indeed any transition system with arithmetical transition relations), the denotation of a Σ^μ_n (resp. Π^μ_n) formula is in Π^1_n (resp. Σ^1_n).

Proof. This is proved by translating the boolean and modal connectives as above (omitting the parameters $\vec{n}$), and then translating the fix-points straight into second-order formulae:

(v′) if $\Phi = \mu Z.\Psi$, then $\phi(m)$ is

$$\forall Z \,.\, (\forall m' \,.\, \psi(m') \Rightarrow Z(m')) \Rightarrow Z(m)$$

which says that m is a member of every pre-fix-point of Ψ;

(vi′) and dually if $\Psi = \nu Z.\Psi$, then $\phi(m)$ is

$$\exists Z \,.\, (\forall m' \,.\, Z(m') \Rightarrow \psi(m')) \wedge Z(m)$$

which says that m is a member of some post-fix-point of Ψ.

Transforming the result to prenex normal form gives the result. (In [Bra91], the author was using the Emerson–Lei definition of alternation, which makes this last calculation easier; with the present definition, slightly more care is required.) □

For lower bounds, [Bra91] outlined a proof that for any Π^1_1 set there is a net and Σ^μ_2 formulae denoting it, by coding the computation of an inductive definition over a Π^1_0 set; later in this paper we give a rather easier proof.

4 An upper bound.

Our main result is that for any reasonable transition system, any mu-calculus formula has a denotation in Δ^1_2, so giving a strong upper bound on the descriptive complexity of sets that one might meet in model-checking. To prove this, we employ a technique of Spector's, which he used to show that positive Σ^1_2 inductively defined sets are in Σ^1_2. Before proceeding to the theorem, we need to introduce another standard coding trick from descriptive set theory.

Let α be an irrational. Define a relation $\leq_\alpha$ on ω by $(m \leq_\alpha n) \Leftrightarrow (\alpha(\langle m, n\rangle) = 1)$. It is a simple exercise to show that '$\leq_\alpha$ is a well-ordering' is expressible by a Π^1_1 formula $WO(\alpha)$. For such an α, we say that α is the code of the ordinal $|\alpha|$ of $\leq_\alpha$; this gives irrational codes for all the countable ordinals. (Note that $\leq_\alpha$ need not be total: if $m \leq_\alpha n$, WO requires that $n \leq_\alpha n$ and $m \leq_\alpha m$, but there may be n such that $n \not\leq_\alpha n$, for instance in the natural codings of the finite ordinals. Given $\alpha \in WO$, the ordinals less than $|\alpha|$ are named by the n such that $n \leq_\alpha n$.)

Theorem 5. Let $\mathscr{T}$ be a transition system with a countable, recursive (or even Δ^1_2) set of (integer codes for) states and set of labels, and with a recursive (or even Δ^1_2) transition relation, and let $\mathscr{V}$ be a valuation assigning recursive (or even Δ^1_2) sets of states to variables. Then for any mu-calculus formula Φ, we have that $\|\Phi\|$ is Δ^1_2.

Proof. We translate $s \in \|\Phi\|$ into arithmetic, starting as we did for nets. Since we are no longer finite-branching, the translations of the modal connectives now need explicit first-order quantifications as in the semantics:

(iii′) if $\Phi = [K]\Psi$, then $\phi(s)$ is $\forall l \in K . \forall s' \in \mathscr{S} . (s \xrightarrow{l} s') \Rightarrow \psi(s')$
and dually for diamond.

The translation for $\Phi = \nu Z.\Psi$ is that of (vi′) above; the difficulty lies in translating $\mu Z.\Psi$ without introducing second-order universal quantifiers. Spector's analysis of least fix-points solves this:

Since $\mathscr{S}$ is countable, the closure ordinal of Φ is countable. Thus (omitting denotational brackets, and writing Φ^ζ for $\mu^\zeta Z.\Psi$) we have

$$s \in \Phi \;\Leftrightarrow\; \exists \kappa < \aleph_1 . \, s \in \bigcup_{\zeta<\kappa} \Phi^\zeta.$$

Now, $\Phi^0, \Phi^1, \ldots$ is a countable sequence of sets of integers (really subsets of $\mathscr{S}$); let κ denote also the ordinal code of κ, and code each Φ^ζ as an irrational α_n, where n is the name of ζ in the coding of ordinal κ given by $\leq_\kappa$. Recall that we can code the infinite sequence $\alpha_0, \alpha_1, \ldots$ by a single irrational α (with arbitrary values for those α_i such that i is not in the domain of $\leq_\kappa$). So we can rewrite $\exists \kappa < \aleph_1 . \ldots$ as $\exists \kappa . \, WO(\kappa) \wedge \ldots$, and if we could compute α, we could rewrite $s \in \bigcup_{\zeta<\kappa} \Phi^\zeta$ as

$$\exists n . \, (n \leq_\kappa n) \wedge ((\alpha)_n(s) = 1). \qquad (*)$$

Now, if we include the definition $\Phi^\zeta = \Psi(\bigcup_{\eta<\zeta} \Phi^\eta)$ of the sets Φ^ζ in our formula, we run into the problem that Ψ would occur negatively as well as positively, thus introducing universal second-order quantification. The trick now is to observe that by the monotonicity of Ψ, instead of using the specific sequence $\Phi^\zeta : \zeta < \kappa$, it suffices to assert the existence of some sequence $X_\zeta : \zeta < \kappa$ of sets that satisfies $X_\zeta \subseteq \Psi(\bigcup_{\eta<\zeta} X_\eta)$, for Φ^ζ is such a sequence, and conversely such a sequence satisfies $X_\zeta \subseteq \Phi^\zeta$, so $(s \in \bigcup_{\zeta<\kappa} X_\zeta) \Rightarrow (s \in \bigcup_{\zeta<\kappa} \Phi^\zeta)$. So we can now further rewrite $s \in \bigcup_{\zeta<\kappa} \Phi^\zeta$ as

$$\begin{aligned} \exists \alpha . \Big[& \forall n . (n \leq_\kappa n) \\ & \quad \Rightarrow \exists \beta . \Big(\forall s'' \in \mathscr{S} . (\beta(s'') = 1) \\ & \qquad\qquad \Leftrightarrow (\exists m . \, m \leq_\kappa n \wedge (\alpha)_m(s'') = 1) \Big) \\ & \qquad \wedge \forall s' \in \mathscr{S} . ((\alpha)_n(s') = 1) \Rightarrow \psi(s)[\beta/Z] \Big] \\ & \wedge (*) \end{aligned}$$

where ψ is the inductively defined translation of Ψ. (The irrational β is the code of $\bigcup_{\eta<\zeta} X_\eta$, owing to the large parenthesized formula.)

Now that we have completed the definition of the translation, we observe that there are no explicit universal second-order quantifiers (and all the explicit existential second-order quantifiers occur positively). Conversion to prenex normal form gives us a block of second order existential quantifiers, followed by a first order formula containing Π^1_1 predicates (in the WO) and Δ^1_2 predicates (the state set $\mathscr{S}$ and the transition relation $\rightarrow$, which occur both negatively and positively in the modal connectives). Therefore ϕ is Σ^1_2, and so $\|\Phi\|$ is Σ^1_2; finally, since the mu-calculus is closed under negation, we have that $\|\Phi\|$ is indeed Δ^1_2. $\square$

One of the STACS referees has remarked that this result is connected to Rabin's result [Rab70] that any formula of SnS is equivalent to a formula $\exists X \,.\, \forall Y \,.\, \chi(X,Y)$ where χ is a weak formula. Indeed, for those who dislike the coding above, this result of Rabin's can also be used to prove the Δ^1_2 upper bound: the modal mu-calculus can be translated into SnS, and since a finite set of integers is codable as a single integer, any weak formula of SnS can be arithmetized as a first-order formula of arithmetic, and so the result follows.

5 Further observations.

We now give a selection of additional results, giving various refinements and applications of the results above.

Firstly, Proposition 4 gave only an upper bound of Π^1_1 for the denotations of Σ^μ_1 formulae (which, recall, is the class of formulae containing only minimal fix-points, with no restrictions on nesting). In fact, we can get a much better upper bound than this for Petri nets, since the continuity lemma used in Theorem 3 in fact also applies to Σ^μ_1 formulae, and so the argument of Theorem 3 goes through to give

Proposition 6. The denotation of an Σ^μ_1 formula on Petri nets is Σ^0_1. $\square$

The introduction of an alternation takes us immediately to second order sets, and there is a simple proof of this:

Proposition 7. For any Π^1_1 set S of integers, there is a net N and an Σ^μ_2 formulae Φ such that $\|\Phi\|$ is the set of markings with all places empty except for one place containing one token, and another place containing n tokens for any $n \in S$.

Proof. We construct a net simulating a register machine as before, but now our machine is one that computes with one integer input x and one irrational input α, using the standard oracle idea briefly mentioned in section 2. So the machine has an input register R_α, and when it requires a further element of α, it transfers control to the state q_{in}. The oracle sets R_α to the next element of α, and changes the state to q_{out}; we model this in the net in the obvious way. Now we add a token generator (as in Theorem 2) for R_α, identifing A_α with q_{in}, and letting the b_α transition transfer control to q_{out}.

Now consider the formula halt(N) from Theorem 1, with the body of the formula extended by the additional conjunct

$$\ldots \wedge ((q_{in} = 1) \Rightarrow \nu Y.[a_\alpha]Y \wedge [b_\alpha]Z)$$

This now means that the original machine halts whatever values are supplied by the oracle, i.e. if $\forall\alpha.\ \phi(\alpha,x)$, where ϕ is the semi-recursive predicate computed by the machine. Since any Π^1_1 set has this form, the result follows. $\square$

Again, we can obtain a tighter upper bound on Σ^μ_2 (for Petri nets) than the trivial one from Proposition 4—and in fact the lower bound we have just proved is also an upper bound.

Proposition 8. For any Petri net and any Σ^μ_2 formula Φ, we have that $\|\Phi\|$ is in Π^1_1.
Proof. The main idea of the proof is this: If ϕ is in Σ^μ_2, then its 'essential structure' is $\mu Y.\ldots(\nu Z.\psi(Y,Z))\ldots$. Thus, for a given value of Y, the inner fix-point has closure ordinal ω. So we can convert the formula to arithmetic by using the easy second-order translation for $\mu Y.$, and the first-order translation (carrying a second-order parameter Y) for the $\nu Z.$, and the result follows. (The slightly messy details are to show that this argument applies to the actual structure given by the definition of Σ^μ_2, but this is not hard—and if one uses the Emerson–Lei alternation classes, it is easier.) $\square$

As an immediate corollary of the last two results, we get the non-collapse of the alternation hierarchy at level 2, previously proven by automata-theoretic arguments [ArN90]:

Corollary 9. There exists a formula Φ in Σ^μ_3 which is not logically equivalent to any formula in Σ^μ_2.
Proof. Take the example net and Φ in Proposition 7 for some Π^1_1-complete set S, and negate it. $\neg\Phi$ is in Π^μ_2, and therefore in Σ^μ_3, but its denotation is not Π^1_1, so it cannot be equivalent to any Σ^μ_2 formula. $\square$

Finally, let us consider the notions of 'reasonable' and 'sufficiently powerful' system that we have used. The transition systems allowed in the main theorem are so powerful that one certainly hopes never to encounter them, so that is reasonable enough. In the arithmetical upper bound results for Petri nets, Theorem 3 and Proposition 6, we made crucial use of the fact that for a marking of a net, there is a predetermined finite set of potential transitions from that marking—this allowed us to avoid any quantification in translating the modal connectives. As long as every state of a system has a recursively bounded finite set of potential transitions, the argument here will go through. If we only require that the system is finite-branching, without a computable bound on the number of branches for a given state, we will have to include quantification in the modalities, and so we lose the easy bounds; however, Proposition 8 still goes through, with a little persuasion. Once we lift the finite-branching restriction, though, we lose everything: it is a simple exercise for the reader to construct a simple infinite branching system and a Σ^{un}_1 formula with an arbitrary Π^1_1 denotation. (Hint: replace the token generator in Proposition 7 by an infinite set of transitions all with the same label.)

On the other hand, our lower bound results are essentially about non-deterministic Turing machines, and then the trick of selecting certain execution paths allowed us to obtain the same results with the strictly less powerful Petri net model. We can in fact obtain the same lower bounds with even smaller classes of model. In the last couple of years, there has been a series of (un)decidability results for model-checking and bisimulation checking on subclasses of Petri nets (see [EsN94] for a survey of these). It turns out that with sufficient ingenuity, one can 'weakly simulate' register machines even with a class of processes (or Petri nets) as small as the so-called Very Basic Parallel Processes, which are nets in which every transition has exactly one input arc and no place has more than one output arc [Esp94]. This construction should also suffice for the token generator trick, although this remains to be checked. (However, for Basic Process Algebra, *alias* context-free processes, in which there is non-determinism but no parallelism, the model-checking problem is known to be decidable.)

6 Conclusions.

We have presented an analysis of the complexity of model-checking infinite systems using a popular temporal logic. By establishing a Δ^1_2 upper bound, we have shown that only relatively simple second-order sets need be considered in proof techniques for such systems (although, of course, many would consider that any second-order set is too complicated!). Furthermore, by establishing lower bounds, we have shown that second-order power is actually needed, even if we restrict attention to rather simple models of systems.

As an interesting corollary, we have re-proved the non-collapse of the alternation hierarchy at level 2. Our aim in future work in this line is to perform a deeper analysis of the classes Σ^μ_i, with the hope of establishing the full non-collapse result that is suspected to be true.

7 Acknowledgements.

The author would like to thank the members of the Concurrency Club at the LFCS, Edinburgh, for helpful discussions, and to thank the STACS referees for their useful comments.

8 References.

[ArN90] A. Arnold and D. Niwinski, Fixed point characterization of Büchi automata on infinite trees. *J. Inf. Process. Cybern.*, **EIK 26**, 451–459 (1990).

[Bra91] J. C. Bradfield, *Verifying Temporal Properties of Systems.* Birkhäuser, Boston, Mass. ISBN 0-8176-3625-0 (1991).

[Bra92] J. C. Bradfield, A proof assistant for symbolic model-checking. *Proc. CAV '92.* LNCS **663** 316–329 (1993).

[Bra95] J. C. Bradfield, On the expressivity of the modal mu-calculus, Technical Report ECS–LFCS–95–338, LFCS, University of Edinburgh (1995). On-line via `http://www.dcs.ed.ac.uk/home/jcb/`.

[BrS92] J. C. Bradfield and C. Stirling, Local model checking for infinite state spaces. *Theoret. Comput. Sci.* **96** 157–174 (1992).

[EmL86] E. A. Emerson and C.-L. Lei, Efficient model checking in fragments of the propositional mu-calculus. *Proc. First IEEE Symp. on Logic in Computer Science* 267–278 (1986).

[Esp94] J. Esparza, On the decidability of model-checking for several μ-calculi and Petri nets, *Proc. CAAP '94*, LNCS **787** 115–129 (1994).

[EsN94] J. Esparza and M. Nielsen, Decidability issues for Petri nets – a survey, *J. Inform. Process. Cybernet.* **30** 143–160 (1994).

[Kai95] R. Kaivola, On modal mu-calculus and Büchi tree automata. *Inf. Proc. Letters* **54** 17–22 (1995).

[Koz83] D. Kozen, Results on the propositional mu-calculus. *Theoret. Comput. Sci.* **27** 333–354 (1983).

[Lar90] K. Larsen, Proof systems for satisfiability in Hennessy–Milner logic with recursion. *Theoret. Comput. Sci.* **72** 265–288 (1990).

[Mos80] Y. N. Moschovakis, *Descriptive set theory.* North-Holland, Amsterdam & New York (1980).

[Niw86] D. Niwiński, On fixed point clones. *Proc. 13th ICALP*, LNCS **226** 464–473 (1986).

[Rab70] M. O. Rabin, Weakly definable relations and special automata, in Y. Bar-Hillel (ed.) *Mathematical Logic and Foundations of Set Theory*, North-Holland, Amsterdam (1970), 1–23.