

These are first-draft lecture notes issued for an optional PG theory course in Edinburgh. The course was concerned with descriptive hierarchies, specifically the arithmetical, analytic, and fixpoint hierarchies. The main aim was to go through the proof of the modal mu-calculus alternation hierarchy; on the way, a quick survey of the basics of effective descriptive set theory is given, together with a few of the more amusing set-theoretic facts. – JCB

1 Introduction

Let us recall some basic facts about computable functions on the natural numbers, which were covered in earlier courses. We know that a (possibly partial) function is computable by a Turing machine iff it can be defined by means of a small set of operators, including in particular primitive recursion and minimalization. Accordingly we call a set of numbers ‘recursive’ if its characteristic function is computable. We know that there exist non-recursive sets and functions; indeed, we know that there are sets that are not even semi-recursive. (Recall that a set is semi-recursive if there is a program which confirms membership, even if it doesn’t necessarily terminate on a non-member. Equivalently, we can list all the members of the set; hence the equivalent term ‘recursively enumerable’, or r.e.)

Given this, we might reasonably ask, what can we say about these ‘hard’ non-r.e. sets? How complicated can they be? And what is the right notion of complicated, anyway?

One approach to answering these questions is to keep our focus on the notation of computation, and start asking questions such as: if somebody gives me some non-recursive set as an oracle, what other sets can I then compute, using a standard Turing machine together with the oracle? If I then take a set I can’t even compute with my oracle, and then add *that* as an oracle, what happens? How far can I go on? Can I order sets in terms of the strength of the oracles needed to compute them? And so on. This approach leads to the extensive area known as degree theory—a degree is effectively a measure of how non-recursive a set is, with respect to other sets. It turns out that the structure of degrees is quite complex. For example, any countable partial order can be embedded in the lattice of degrees, and there is an antichain of cardinality 2^{\aleph_0} .

However, that is not the approach I want to discuss. Rather, let’s draw back a little from computation: after all, if we can’t compute a set, what’s the point in discussing just how badly we can’t compute it? How did we get this nasty set, anyway? Well, probably someone presented us with a problem which defined the set in some way. For example, we were asked “is x a zero of the function f ?”. Or, we might have been asked “does process P satisfy the modal mu-calculus Φ ?” In other words, we have a *description* of the set, as “the zeros of f ” or “the states satisfying Φ ”.

This suggests the question, what can we say about how hard it is to describe sets? Now usually we describe things using predicate calculus (or rather, some English that we could in principle turn into predicate calculus), so the complexity of the describing formula is an obvious measure. But what complexity? One obvious, and easy to work with, definition is the number of alternations between existential and universal quantifiers . . .

This course is intended to give overviews of the ideas, and avoid the tricky technical details. However, if you need to use this stuff, or are just interested, you have to read the details some time. There are several books on recursion theory, and all of them contain the details of the first part of this course (the arithmetical and analytical hierarchies). The main library has a copy of

the now out-of-print classic *Theory of recursive functions and effective computability* by Rogers. A newer book is by Shoenfield, called *Recursion Theory*. Moschovakis' book *Descriptive Set Theory* is a wonderful text covering both the effective theory we're looking at, and the classical theory, and many advanced topics. However, it assumes a basic knowledge of topology and analysis (only very basic, though).

2 The arithmetical sets.

Let us fix the world we're describing, and the language we're using to describe it. For this section, we take \mathbb{N} as the structure, and we describe subsets of \mathbb{N} by means of formulae in the first-order language of recursive arithmetic. That is, we have first-order predicate logic, together with constants for all the *recursive relations* on \mathbb{N}^k , for any k . Note that functions are not built in: a function f is described in terms of its graph $G^f = \{ (x, y) \mid y = f(x) \}$. However, if f is a *total* recursive function, then its graph is recursive, so we may as well add constants for recursive functions as well. In particular, it's essential to have functions that code a finite sequence x_1, \dots, x_k (for any k) of integers into a single integer $\langle x_1, \dots, x_k \rangle$, and corresponding functions $\text{len}(u)$ and $(u)_i$ to give the length of a coded sequence and extract its i th element. (Boring exercise: define such functions.)

We shall assume that we may freely use the standard laws of predicate logic; in particular, the De Morgan laws can be used to push negation inwards.

Now we define a hierarchy of formulae according to the alternation of quantifiers, thus:

Definition 1 If a formula ϕ contains no quantifiers, it is Σ_0^0 and also Π_0^0 . Otherwise, for $n \geq 1$:

- If ϕ is Σ_{n-1}^0 or Π_{n-1}^0 , it is both Σ_n^0 and Π_n^0
- If ϕ_1 and ϕ_2 are Σ_n^0 (resp. Π_n^0), then so are $\phi_1 \vee \phi_2$, $\phi_1 \wedge \phi_2$ and $\exists x. \phi_1$ (resp. $\forall x. \phi_1$).
- If ϕ is Σ_n^0 (resp. Π_n^0), then $\forall x. \phi$ (resp. $\exists x. \phi$) is Π_{n+1}^0 (resp. Σ_{n+1}^0).

In other words, a Σ_n^0 formula has n blocks of quantifiers, starting with an \exists block. Note that the classes Σ_0^0 and Π_0^0 are often not defined in presentations of the subject; and since they don't really fit very well, unless otherwise stated all future references to Σ_n^0 etc. should be taken to assume $n \geq 1$.

Exercise 2

- (1) Verify that ϕ is Σ_n^0 iff $\neg\phi$ is Π_n^0 .
- (2) Why is it reasonable that we don't consider $\exists x. \exists y. \phi$ to be any more complex than $\exists x. \phi$?
- (3) The prenex normal form theorem for first-order arithmetic says that any Σ_n^0 formula ϕ is equivalent to a formula of the form $\exists \vec{x}_1. \forall \vec{x}_2. \dots (\exists/\forall) \vec{x}_n. \psi$, where the last quantifier is \exists or \forall according as n is odd or even, and ψ is quantifier-free. Prove this. For this reason, the definitions of Σ_n^0 usually assume prenex normal form, and the closure of Σ_n^0 under \vee, \wedge, \exists is proven as a theorem rather than stated in the definition. We choose to put it in the definition because we shall later work with a logic, the modal mu-calculus, that doesn't have a prenex normal form.
- (4) Verify that a relation $T(\vec{x})$ is semi-recursive iff there is a recursive relation $R(y, \vec{x})$ such that $T(\vec{x}) \Leftrightarrow \exists y. R(y, \vec{x})$.

The last exercise says that a relation is semi-recursive iff it is definable by a Σ_1^0 formula. It should perhaps be noted here that this is not true in all the spaces on which we might define recursive functions. It is possible to develop the theory in a somewhat more general setting, and in that case the last exercise does not necessarily hold. One then has to *define* Σ_1^0 to be the semi-recursive relations, and Π_1^0 to be their complements. (For example, the exercise fails in \mathbb{R} .)

Now we can define a hierarchy, the *arithmetical hierarchy*, of definable sets and relations:

Definition 3 A relation $T(\vec{x})$ is Σ_n^0 (Π_n^0) if there is a Σ_n^0 (Π_n^0) formula $\phi(\vec{x})$ such that $T(\vec{x}) \Leftrightarrow \phi(\vec{x})$.

Warning: the definition says that ϕ exists; it doesn't say that we know what it is. For example, consider the set T of integers defined by

$$T(x) \Leftrightarrow (x = 1) \wedge \text{the Riemann hypothesis is true.}$$

The Riemann hypothesis is not even expressible as a first-order sentence about \mathbb{N} , and it's certainly very difficult, so we might reasonably consider T to be a very difficult set. However, either it is the case that $T = \emptyset$, or it is the case that $T = \{1\}$, and both of these sets are as about as simple as possible. The point is that if we were clever enough, we could describe T with a Σ_0^0 formula; whereas, as we shall see, there are sets that simply can't be described without using umpteen quantifiers. (Incidentally, you might like to consider the nature of the set defined by “ $x = 1$ and the continuum hypothesis holds”.)

It is also useful to have a notation for sets that are both Σ and Π :


Definition 4 A relation is Δ_n^0 if it is both Σ_n^0 and Π_n^0 .

Note that this is a property of relations, not of formulae. Note also that $\Delta_1^0 = \Delta_0^0 = \text{recursive}$.

Exercise 5 The Σ_n^0 classes are closed under \exists but not \forall , and the Π_n^0 classes are closed under \forall but not \exists . Thus the Δ_n^0 classes are not closed under quantification. However, there is a restricted form of quantification, *bounded quantification*, which is often useful: the bounded universal quantification of $\phi(x)$ is $\forall x < y. \phi$ (note that y is a free variable), and similarly for existential. It's trivial that Π_n^0 is closed under bounded universal quantification, since $(\forall x < y. \phi) \Leftrightarrow (\forall x. (x < y) \Rightarrow \phi)$.

Show that Σ_n^0 is closed under bounded universal quantification. (Hint: use the sequence coding functions to represent a suitable Skolem function as a single integer.) Thus all the classes are closed under bounded quantification.

We have the following diagram, which is (in our formulation) trivial to obtain:



The next task is to show that all these inclusions are indeed strict. We already know that there are semi-recursive relations that are not recursive; the strictness of remaining inclusions are shown by a similar technique of uniformization and diagonalization. To do these proofs properly requires tedious attention to detail in the form of coding techniques; we'll not worry about this, and so the word ‘proof’ should be fairly loosely interpreted in these notes. The real proofs can be found in the standard texts.

The main result is the parameter theorem for Σ_n^0 :

Theorem 6 *There are universal Σ_n^0 formulae: that is, for each $n \geq 1$ and $k \geq 1$, there is a Σ_n^0 $(k + 1)$ -ary relation U_n^k such that for each Σ_n^0 k -ary relation R there is an integer r such that $R(\vec{x}) \Leftrightarrow U_n^k(r, \vec{x})$; and dually there are universal Π_n^0 formulae V_n^k .*

Proof. By induction on n . The base case is Σ_1^0 ; but this is Kleene's Parameter Theorem, so we already know it (fortunately, since it's where the hard work happens).

So suppose we have the result for Σ_n^0 and Π_n^0 , and let R be a Σ_{n+1}^0 k -ary relation. Then $R(\vec{x}) \Leftrightarrow \exists y. S(y, \vec{x})$ where S is Π_n^0 . By induction, there is an s such that $S(y, \vec{x}) \Leftrightarrow V_n^{k+1}(s, y, \vec{x})$ where V_n^{k+1} is Π_n^0 . So define $U_{n+1}^k(r, \vec{x}) \Leftrightarrow \exists y. V_n^{k+1}(r, y, \vec{x})$ and we're done. Negating U_{n+1}^k gives us V_{n+1}^k . \square

Corollary 7 For $n \geq 1$ and $k \geq 1$ there is a Σ_n^0 k -ary relation that is not Π_n^0 .

Proof. Consider $U = U_n^k$. Define the k -ary relation $W(y, \vec{x}) = \neg U(y, y, \vec{x})$. Since U is Σ_n^0 , W is Π_n^0 . Suppose that W is also Σ_n^0 . Then it has an index w such that $W(y, \vec{x}) \Leftrightarrow U(w, y, \vec{x})$. Now consider $W(w, \vec{x})$. We have $W(w, \vec{x}) \Leftrightarrow \neg U(w, w, \vec{x}) \Leftrightarrow \neg W(w, \vec{x})$, contradiction. So W can't be in Σ_n^0 , and $\neg W$ can't be in Π_n^0 . \square

This result establishes that all the inclusions in the hierarchy diagram are strict (exercise: why?).

Exercise 8 Define a relation on \mathbb{N} that is not Σ_n^0 for any n . (Hint: take n as one of the arguments of the relation.)

Finally, let's reconsider the language we're using, in particular the stock of constants. It may seem unduly generous to allow all recursive relations as constants, particularly since we don't have any way of telling whether a given relation (presented as a Turing machine, or recursive definition) *is* recursive! This isn't much of a problem, since the constants we actually want to use are always relations that are very easily shown to be recursive. However, the following amazing theorem, which resolved (negatively) Hilbert's Tenth Problem, tells us that we don't need to be so profligate with constants.

Theorem 9 (Matijacevič' Theorem) If $R(\vec{x})$ is a semi-recursive relation, it can be defined by a formula of the form $\exists \vec{y}. p(\vec{y}, \vec{x}) = q(\vec{y}, \vec{x})$, where p and q are polynomials over \mathbb{N} .

So it suffices to take the constants $0, 1, +, \cdot$, and everything remains unchanged (apart from the definition of Σ_0^0 , which doesn't really count).

3 Second-order arithmetic.

3.1 Semi-recursive relations on $\mathbb{N} \rightarrow \mathbb{N}$.

Now that we've dealt with first-order arithmetic, the obvious next step is second-order. One might ask whether this is necessary—do we ever see non-first-order sets in real life? The answer is 'yes', for two reasons. Firstly, it turns out that such computer-science stalwarts as fix-points and games take us quickly into second-order territory; and secondly, most of the real world is concerned with \mathbb{R} , not \mathbb{N} , and first-order statements about \mathbb{R} turn into second-order statements about \mathbb{N} .

What do we mean by second-order? We mean that as well as quantifying over \mathbb{N} , we can quantify over functions $\mathbb{N} \rightarrow \mathbb{N}$. Functions are slightly abstract objects, and it's helpful to think of a function $\alpha: \mathbb{N} \mapsto \mathbb{N}$ as an infinite sequence $\alpha(0), \alpha(1), \dots$ of numbers. Of course, if we're going to quantify over functions, all our definitions of 'recursive' and so on need to be revised to apply to relations that may be over functions as well as numbers; this means we need to know what a (semi-)recursive set of functions is, which may not be entirely obvious. However, if we follow our computational intuition, everything works out.

So, suppose we have a unary relation $R(\alpha)$ over functions—what should it mean for R to be semi-recursive? (As I noted earlier, semi-recursive-ness is the primitive notion, not recursive-ness.) It should mean that we can write a program which will determine membership of R , in the sense that if $R(\alpha)$, then the program should terminate with the answer 'yes' when given input α . This raises the question of how a function can be given as input to a program. There are two ways of thinking about the natural answer to this: either we can think of the machine being fed the infinite sequence $\alpha(0), \alpha(1), \dots$ of function values, which it can read as it desires; or we can think of the machine being given α as a black box or oracle, so that whenever it wants to know $\alpha(n)$, it just asks the oracle. (Note that we are absolutely not placing any restrictions on how hard α is to compute; we're not even trying to compute it, it's just there.) Given this, it's clear that the only way the program can terminate is if it only needs a finite number of function values in order to answer the question of whether $\alpha \in R$.

This suggests the following definition:

Definition 10 Let \mathcal{N} denote the space $\mathbb{N} \rightarrow \mathbb{N}$, viewed as the space of infinite sequences. Let $r_i, i \in \mathbb{N}$ be some enumeration of the set of finite sequences of numbers, and define the i th basic recursive set to be

$$B'_i = \{ \alpha \mid \forall j \leq \text{len } r_i. \alpha(j) = (r_i)_j \}$$

and say that a relation $R(\alpha)$ is semi-recursive if it is of the form

$$R = \bigcup_{i \in \mathbb{N}} B'_{\epsilon(i)}$$

for some recursive function ϵ .

We sometimes use the word *irrational* to denote a member of \mathcal{N} (see the following section for the reason).

This definition coincides with the intuition above, since to determine that α is in R , we have check it against a recursively enumerated list of possible finite prefixes, and this will terminate

if the answer is ‘yes’. Note that the function ϵ is just a function $\mathbb{N} \rightarrow \mathbb{N}$, so it’s a member of \mathcal{N} .

This definition works fine for \mathcal{N} , and there’s an obvious generalization for relations involving any (finite) number of arguments from \mathbb{N} and \mathcal{N} . However, it is actually a simplified version of a more general definition. Understanding the general version is helpful, but requires a little topology; I’ll therefore sketch it in the following section, but feel free to omit this.

3.2 Recursiveness in general spaces.

The general theory is developed for *perfect Polish spaces*. That is, a metric space that is complete (every Cauchy sequence converges to a limit) and separable (contains a countable dense set) with no isolated points. Most reasonable spaces are such things. In particular, \mathcal{N} (called *Baire space*) is a perfect Polish space if we give it the following metric d , which is a metric for the natural product topology on $\mathcal{N} = \mathbb{N}^{\mathbb{N}}$:

$$d(\alpha, \beta) = \frac{1}{1 + \min\{n \mid \alpha(n) \neq \beta(n)\}}$$

(where, of course, we take $1/\infty = 0!$). \mathbb{R} is also such a space; and so is Cantor space, the space of functions $2 \rightarrow \mathbb{N}$, with the same metric as \mathcal{N} . Any finite product of perfect Polish spaces is one also.

Given a perfect Polish space \mathcal{X} , there’s an obvious neighbourhood basis, namely the open balls $B(x, m/(n+1)) = \{y \mid d(x, y) < m/(n+1)\}$ for $x \in \mathcal{X}$ and $m, n \in \mathbb{N}$. Unfortunately, although there are only countably many rationals to put in the second place, there are still uncountably many x to put in the first place. However, all interesting perfect Polish spaces admit a *recursive presentation*: that is, an infinite sequence r_0, r_1, \dots of points, which is dense in \mathcal{X} and such that the following relations are recursive:

$$\begin{aligned} P(i, j, m, n) &\Leftrightarrow d(r_i, r_j) \leq m/(n+1) \\ Q(i, j, m, n) &\Leftrightarrow d(r_i, r_j) < m/(n+1); \end{aligned}$$

that is, given any two of the points, and some rational distance D , we can effectively tell whether they’re more or less than D apart. In the case of \mathcal{N} , the set of ultimately zero sequences forms a recursive presentation (that’s what we were really doing above).

Exercise 11 Give a recursive presentation of \mathbb{R} .

If \mathcal{X} has a recursive presentation, then we have a recursively enumerable neighbourhood basis $N(i), i \in \mathbb{N}$ for \mathcal{X} , by taking some convenient enumeration of the balls $B(r_k, m/(n+1))$. We can now define

Definition 12 A relation R on \mathcal{X} is semi-recursive if it has the form $\bigcup_{i \in \mathbb{N}} N(\epsilon(i))$ for some recursive $\epsilon: \mathbb{N} \rightarrow \mathbb{N}$.

Although this is not quite the same as the simplified definition we had before, it’s morally equivalent for \mathcal{N} . However, this full definition is required for spaces such as \mathbb{R} .

Exercise 13 If you’re wondering what is so different about \mathbb{R} , consider this. The main difference is that given $x, y \in \mathbb{R}$, we need to know all the digits of both x and y in order to calculate $d(x, y) = |x - y|$, whereas with \mathcal{N} we only need a finite number of digits, unless

$x = y$. Topologically, \mathbb{R} is a connected space. \mathcal{N} , however, is totally disconnected—and so the basic open sets B are in fact also closed.

Now in general, a relation R is *defined* to be recursive if both R and its complement are semi-recursive. So, using the above definition, if a relation R on \mathbb{R} is recursive, then it is a union of open sets, and also an intersection of closed sets. What does this say about R , and why does it show that Exercise 2(4) fails on \mathbb{R} ?

We may remark here that \mathcal{N} is actually homeomorphic to the set of irrational numbers, viewed as a subspace of \mathbb{R} ; and hence the elements of \mathcal{N} are often called *irrationals*.

Having dealt with perfect Polish spaces, it's easy to throw in some copies of \mathbb{N} ; so the general spaces on which descriptive set theory is done are products of the form $\mathbb{N}^k \times \mathcal{X}_1 \times \dots \times \mathcal{X}_m$, where each \mathcal{X}_i is some basic space such as \mathcal{N} or \mathbb{R} . Given a space of this form, one can imagine projecting it down one of the copies of \mathbb{N} ; this corresponds to a first-order existential quantification. Thus one can build up complex relations by applying projection and complementation to existing relations—and of course one can project along one of the \mathcal{X}_i , a second-order projection. So for the classical descriptive set theorist, the logical description we're using is just a description of the underlying operations on our spaces. However, it makes life a lot easier for us to take the logic as primary, so we'll stick with that.

That ends (for the moment) our brief diversion into the general setting—we now return to the special case of \mathbb{N} and \mathcal{N} .

3.3 The analytical hierarchy.

Everything we did for the naturals extends to relations on \mathbb{N} and \mathcal{N} , as far as the arithmetical hierarchy goes. So now we start applying second order quantification. From now on, I'll only state definitions for Σ , and assume the dual definition for Π . So, our language is now second-order arithmetic; and we shall have the convention that α, β, \dots range over \mathcal{N} , and most roman letters range over \mathbb{N} .

Definition 14 If a formula ϕ is Σ_1^0 , then it is Σ_0^1 . Otherwise, for $n \geq 1$,

- if ϕ is Σ_{n-1}^1 or Π_{n-1}^1 , it is both Σ_n^1 and Π_n^1 ;
- if ϕ_1 and ϕ_2 are Σ_n^1 , so are $\phi_1 \vee \phi_2$, $\phi_1 \wedge \phi_2$, $\exists x. \phi_1$, $\forall x. \phi_1$, and $\exists \alpha. \phi_1$;
- if ϕ is Σ_n^1 , then $\forall \alpha. \phi$ is Π_{n+1}^1 .

Definition 15 A relation R on \mathbb{N} and \mathcal{N} is Σ_n^1 if it is definable by a Σ_n^1 formula. It is Δ_n^1 if it is both Σ_n^1 and Π_n^1 .

Exercise 16

- (1) Show that it is possible to code an *infinite* sequence $\alpha_0, \alpha_1, \dots$ of irrationals into a single irrational $\langle \alpha_0, \alpha_1, \dots \rangle$ by recursive coding. (Hint: it's just like the standard mapping of $\mathbb{N} \times \mathbb{N}$ into \mathbb{N} .)
- (2) As with first-order arithmetic, there is a prenex normal form theorem: any Σ_n^1 formula is equivalent to one of the form $\exists \vec{\alpha}_1. \forall \vec{\alpha}_2. \dots (\exists/\forall) \vec{\alpha}_n. (\forall/\exists) \vec{x}. \phi$ (for n odd/even), where ϕ is quantifier-free; and so the hierarchies are customarily defined assuming normal form, and the closure properties proved. Prove the normal form theorem. (Hint: you need to push quantifiers around. The tricky part is showing that $\forall x$ can be pushed inside $\exists \alpha$. Use part (1) to do this. Then to get the induction going, you need to show that any first-order formula can be put into the form $\exists \alpha. \forall x. \phi$, by Skolemizing and pushing $\forall x$'s inside. (This incidentally shows that $\bigcup_n \Sigma_n^0 \subseteq \Delta_1^1$.)

(3) Show that if $S(\alpha, \vec{x})$ is Σ_1^0 and R is defined by $R(\vec{x}) \Leftrightarrow \exists \alpha. S(\alpha, \vec{x})$, then R is also Σ_1^0 .

There is a useful fact about both the arithmetical and analytical hierarchies, which you probably assume without thinking about it, but which actually requires proof. The term *Kleene pointclass* refers to any of the $(\Sigma/\Pi/\Delta)_n^{(0/1)}$ (the term ‘pointclass’ is from the jargon of descriptive set theory, and means ‘set of relations’). We’ll state it for the record, but shan’t prove it. The enthusiastic reader may wish to sketch the proof.

Theorem 17 *The Kleene pointclasses are closed under recursive substitution. That is, if $R(x, \dots)$ is a member of a Kleene pointclass, and f is a recursive function, then $R(f(x), \dots)$ is also in the Kleene pointclass.*

Exercise 18 This is also true when x is an irrational: what is a recursive function $\mathcal{N} \rightarrow \mathcal{N}$?

We have the obvious inclusion hierarchy, as for the arithmetical hierarchy. It remains to prove that the inclusions are strict.

Theorem 19 *There are universal Σ_n^1 relations, and hence the analytical hierarchy is strict.*

Proof. Exactly analogous to the proof for the arithmetical hierarchy.

4 The hyperarithmetical hierarchy.

In this section, I'd like to say a few things about what happens between the top of arithmetical hierarchy and the bottom of the analytical hierarchy. We've already seen (if you did Exercise 8) an example of a non-arithmetical set; and we've seen artificial examples of arbitrarily hard analytical sets. One might wonder how big the gap between arithmetical and, say, Δ_1^1 , really is. It turns out to be quite large, in some senses—but quite small in another!

If you followed the hint in Exercise 8, you probably produced a set looking something like

$$T(n, w, x) \Leftrightarrow U_n(w, x)$$

where U_n is the universal Σ_n^0 relation. This set is obviously not arithmetical, but it's not very unarithmetical—it's just a union or disjunction of arithmetical sets, since each U_n is arithmetical. Moreover, it's a very well-behaved union: it's a union of countably many sets, and we know effectively for each i what the i th set in the union is. There's an obvious similarity between such a union and the operation of existential quantification— $\exists x. P(x, y) \Leftrightarrow y \in \bigcup_{i \in \mathbb{N}} \{y \mid P(i, y)\}$ —and this might suggest to us that the arithmetical hierarchy could be extended into the transfinite. Of course, to keep our notion of effectiveness, we must take well-behaved, i.e. recursive, unions, rather than arbitrary unions. This leads us to the following definition:

Definition 20 A relation T is Σ_ζ^0 for a recursive ordinal ζ if $T = \bigcup_{i \in \mathbb{N}} R_{f(i)}$, where the R_j enumerate the relations in $\bigcup_{\xi < \zeta} \Pi_\xi^0$, and f is a recursive function.

Several things in this definition require comment. Firstly, we've defined Σ_ζ^0 directly on relations, rather than on formulae, contrary to our previous policy. There's no difficulty about doing it logically, but it requires us to extend our previous definitions to infinitary first-order logic, which isn't worth the effort for the small use we'll make of this definition. Secondly, we've said that ζ is a *recursive ordinal*—what is that, and why do we require it? Well, we require it to keep things effective: we can't hope to have a recursive enumeration of the $\bigcup_{\xi < \zeta} \Pi_\xi^0$ relations unless we can recursively enumerate the ordinals below ζ , and we can only do that if ζ is recursive: that is, there is a recursively coded notation for the ordinals less than ζ , such that we can recursively determine for any two codes which is less than which.

Exercise 21 Convince yourself that if we restrict to the finite ordinals, the above definition agrees with our existing definition of Σ_n^0 relations. (You'll need the parameter theorem.)

We can deal formally with recursive ordinals as follows. Let \prec be a binary relation on \mathbb{N} which is a well-order. (Write down the formal expression of 'being a well-order'.) The order type of \prec is some (countable) ordinal ζ . Then \prec gives us integer notations for all the ordinals less than ζ — ξ is denoted by the ξ th number in the well-ordering. If \prec is recursive, then we can effectively determine $\xi_1 < \xi_2$ by comparing their notations. Thus we can certainly enumerate the ordinals less than any given ξ , provided we know the code of ξ . However, it's not at all obvious how to find the code of ξ —indeed, it's not obvious how to find the code of 0! Not only is it not obvious, it's not possible in general, which may seem a bit of a drawback. An alternative approach is instead to define directly ordinal notation systems—for example, Cantor normal form gives notations for all the ordinals below ϵ_0 , and one can go much further. In fact, one can give a sensible notation system for all the recursive ordinals; but then the \prec relation is only a semi-recursive partial well-order, not a recursive total well-order. However, that's enough for the above definition.

The smallest non-recursive ordinal is called ω_1^{CK} (Church–Kleene ω_1). It is the smallest ordinal that is not the order type of a recursive well-ordering of \mathbb{N} , and has many other characterizations. Our definition of Σ_ζ^0 above only really makes sense for the ordinals below ω_1^{CK} , and so the hierarchy stops naturally there. The relations in $\bigcup_{\zeta < \omega_1^{\text{CK}}} \Sigma_\zeta^0$ are called the *hyperarithmetical* relations. One of the most important results in descriptive set theory is Kleene’s theorem

Theorem 22 *A relation R is hyperarithmetical iff it is Δ_1^1 .*

Thus the extension of the arithmetical hierarchy into the ‘recursively transfinite’ takes us exactly up to the bottom of the analytical hierarchy. The proof of this theorem is quite hard, and involves a normal form for the Π_1^1 relations in terms of well-orderings.

5 Notes on the classical theory.

5.1 Basic definitions and results.

The theory we’ve been doing so far appears to stand by itself. However, it can be seen as an effective version of an older theory, classical descriptive set theory. This theory started when people wanted to analyse the descriptive complexity of sets of real numbers, rather than integers. The idea was that the open sets are the simplest sets of reals, and then one forms more complex sets by operations such as countable union or existential quantification, complementation, and so on. It turns out that you don’t have to go very far into the classical equivalent of the analytical hierarchy, for set-theoretic hypotheses to force their attentions upon you.

We’ll now glance very quickly at some features of the classical theory. We’ll consider relations on finite products of \mathbb{N} and \mathcal{N} (or in general, \mathbb{N} and perfect Polish spaces). \mathcal{N} comes equipped with a topology (see section 3.2).

Definition 23 A relation R is Σ_1^0 if it is open (in the appropriate product topology). R is Π_n^0 if $\neg R$ is Σ_n^0 . If T on $\mathbb{N} \times \mathcal{X}$ is Π_n^0 , then $R = \{x \in \mathcal{X} \mid \exists i \in \mathbb{N}. (i, x) \in T\}$ is Σ_{n+1}^0 . If R is both Σ_n^0 and Π_n^0 , it is Δ_n^0 .

This is the traditional definition. Equivalently, we could define these *bold-face* classes in terms of formulae, just as we did for the light-face classes. For general reasons, we start with Σ_1^0 being the open sets, the equivalent of semi-recursive sets in the effective theory; in fact, for \mathcal{N} we can if we wish start with Σ_0^0 being the clopen sets, but this doesn’t work in general (see Exercise 13).

The bold-face versions of the parameter and hierarchy theorems are true, and easier than the light-face versions: to prove the Kleene parameter theorem (there is a universal Σ_1^0 relation), one can’t avoid all the work of coding recursive definitions, whereas to show there’s a universal Σ_1^0 relation one only needs some manipulation of basic neighbourhoods and Cauchy sequences.

Note, incidentally, that *any* set of naturals is Σ_1^0 , since the only topology on \mathbb{N} is the discrete topology in which all sets are open (and therefore closed). This is why the classical theory has nothing useful to say about sets of naturals.

Definition 24 A relation R is Σ_1^1 if it is Σ_1^0 . R is Π_n^1 if $\neg R$ is Σ_n^1 . If T on $\mathcal{N} \times \mathcal{X}$ is Π_n^1 , then $R = \{x \in \mathcal{X} \mid \exists \alpha \in \mathcal{N}. (\alpha, x) \in T\}$ is Σ_{n+1}^1 . If R is both Σ_n^1 and Π_n^1 , it is Δ_n^1 .

Again, this is the traditional definition, and again, we could equally well cast it in logical terms. And the hierarchy theorem is true.

There is actually a tight relationship between the bold-face and light-face classes:

Theorem 25 *If Γ is any of $(\Sigma/\Pi/\Delta)_n^{(0/1)}$, then a relation R on \mathcal{X} is Γ iff there is a Γ relation S on $\mathcal{N} \times \mathcal{X}$ and an irrational ϵ such that $R(x) \Leftrightarrow S(\epsilon, x)$.*

Proof. Any open set is a countable union of basic neighbourhoods, and we have a countable basis $N(i)$ such that each $N(i)$ is semi-recursive (see section 3.2), so if R is open, i.e. Σ_1^0 , it satisfies $R(x) \Leftrightarrow \exists i. N(\epsilon(i))(x)$ for some irrational ϵ , and this right-hand side is Σ_1^0 . Then a simple induction extends the result to the other Σ/Π classes, and a small coding trick deals with Δ .

We can also extend the Σ_n^0 classes into the transfinite. Since we're not worrying about effectiveness, there's no need to restrict ourselves to recursive unions; we can take arbitrary countable unions.

Definition 26 A relation T is Σ_ζ^0 for a countable ordinal ζ if $T = \bigcup_{i \in \mathbb{N}} R_i$, where each R_i is in $\bigcup_{\xi < \zeta} \Pi_\xi^0$.

This hierarchy extends up to ω_1 , the first uncountable ordinal. It is called the *Borel hierarchy*, and sets in the hierarchy are the *Borel sets*. (The Σ_n^0 don't have a name of their own; they're just the 'Borel pointclasses of finite order'. The Σ_n^1 are called the *Lusin* or *projective* pointclasses. Warning: the Σ_1^1 sets are traditionally called *analytic* sets, not to be confused with the analytical sets!) The classical equivalent of Kleene's theorem is Suslin's theorem

Theorem 27 *A relation is Borel iff it is Δ_1^1 .*

There is a huge body of work analysing the structure of these classes. However, here I just want to give a few examples of some of the remarkable connections between descriptive set theory and the wilder reaches of set theory.

5.2 The impact of set theory on descriptive set theory.

Throughout these notes, we've been implicitly assuming that we're working in *ZFC*, Zermelo–Fraenkel set theory with the full Axiom of Choice (*AC*). (Recall that *AC* asserts that given any set $\{S_x \mid x \in X\}$ of non-empty sets indexed by the set X , there exists a function $f: X \rightarrow \bigcup_{x \in X} S_x$ such that $f(x) \in S_x$.) As is well known, there are other axioms besides choice which one might (or might not) believe should be true, and which are independent of the other axioms of *ZF*.

The first such axiom is *CH*, the *Continuum Hypothesis*, and its *Generalized* version *GCH*. *CH* states that $\aleph_1 = 2^{\aleph_0}$; in other words, every set of reals (or irrationals) is finite, countably infinite, or has the same cardinality as \mathbb{R} (or \mathcal{N}). (\aleph_0 denotes the first infinite cardinal (ω), and so on. \aleph_1 was also called ω_1 above, when we were using it in its role as an ordinal.) *GCH* states the general version $\aleph_{\lambda+1} = 2^{\aleph_\lambda}$. Gödel and Cohen proved that *GCH* is independent of *ZFC* (although *GCH* implies *AC*). In fact, *CH* can be false in a very strong sense, as we'll remark below. However, if we restrict our attention to reasonably simple sets (and in the mathematics that is applied to the real world, most sets are simple), we don't have to worry so much:

Theorem 28 **CH* is true of Σ_1^1 sets.*

In other words, every infinite Σ_1^1 set is either countable or has the cardinality of \mathcal{N} . There are plenty of other curious results about cardinals and the bottom end of the projective hierarchy: for example:

Theorem 29

- Every Σ_1^1 set is a union (or intersection) of \aleph_1 Borel sets.
- (easy) Any Σ_1^1 well-ordering of \mathcal{N} is countable.
- (not easy) Any Σ_2^1 well-ordering has order type less than \aleph_2 ; and so if *CH* is false, there is no Σ_2^1 well-ordering of \mathbb{R} or \mathcal{N} .

Another property that one might want to be true of the world is that all sets of reals are *Lebesgue measurable*. Roughly, this means that every set of reals has a sensible ‘length’, so one can do integration etc. For the sets and functions one meets in practice, this isn’t a problem, but it’s not obvious (or true) that all sets of reals are Lebesgue measurable. Again, if we stick to simple sets we’re OK. Let *LM* mean ‘sets of reals are Lebesgue measurable’.

Theorem 30 *LM holds for Σ_1^1 sets (and indeed for countable unions and intersections of Σ_1^1 and Π_1^1 sets).*

However, it is consistent with *ZFC* that there are Δ_2^1 sets that are not measurable. Full *LM* is not consistent with *ZFC*, since full choice always allows the construction of non-measurable sets; but it is probably consistent with *ZF + DC*, where *DC* is the Axiom of Dependent Choice, which states that one can make a countable sequence of choices, each depending on the previous choices. This is a weak form of choice, but it does provide enough choice for the theory of Lebesgue measure to make sense. The ‘probably’ in the last but one sentence needs explanation, and brings us to one of the most surprising links, between Lebesgue measurability and the existence of large cardinals.

An infinite cardinal κ is *regular* if it is not the sum of less than κ cardinals each less than κ . For example, $\aleph_0, \aleph_1, \dots$ are regular, but \aleph_ω is not, since $\aleph_\omega = \sum_{i < \omega} \aleph_i$. An uncountable cardinal κ is *weakly inaccessible* if it is regular and is a limit cardinal, i.e. is \aleph_λ for some limit ordinal λ (note that by regularity, it must then be the case that $\lambda = \kappa$). Such a cardinal, if it exists, is very large: it’s impossible to reach it ‘from below’ by any form of summation or limit. If moreover $2^\lambda < \kappa$ for all $\lambda < \kappa$, then κ is (*strongly*) *inaccessible*—it can’t be reached from below even with the help of the powerset operation. Let *IC* denote the assertion that an inaccessible cardinal exists. *IC* implies that *ZF* is consistent (because an inaccessible cardinal is a model for *ZF*), and hence it cannot be proven from *ZF*, and further cannot be shown to be consistent with *ZF*; this makes it different in kind from properties such as *GCH*. Thus it remains possible that *ZF* can prove that there is no inaccessible cardinal, though everybody would be astonished if this were ever done.

How does this connect with measurability? Solovay proved that *ZF + DC + LM* is consistent, but to do this, he had to assume *IC* (why, would be too complicated to go into). The surprise came when Shelah proved that this is necessary: the consistency of *ZF + DC + LM* implies the consistency of *ZF + IC*. (Note carefully, by the way, that we are *not* saying that *IC* implies *LM* or vice versa; we are saying that their consistency is equivalent.)

IC is only the smallest of many ‘large cardinal’ axioms. One of the largest is *MC*, the assertion that a measurable cardinal exists. The definition of a measurable cardinal is a bit complicated, but if one exists, it must, in the presence of choice, be really huge: much bigger

than just inaccessible. (Without choice, even \aleph_1 could be measurable!) However, if MC is true, then every Σ_2^1 set is measurable.

It was mentioned above that CH can fail in very strong ways. What this means is that it is consistent with $ZF(C)$ that $2^{\aleph_0} = \aleph_\beta$ for any ‘reasonable’ β . In particular, it is consistent that there is a weak inaccessible less than 2^{\aleph_0} (assuming that the existence of weak inaccessibles is consistent, that is)! Indeed, if you don’t believe that $2^{\aleph_0} = \aleph_1$, there’s a reasonable argument for saying that the first weak inaccessible is the next plausible value to take.

5.3 Games.

Finally, a few remarks on how games tie up with descriptive set theory and large cardinal axioms.

A *Gale–Stewart* game is specified by a set $A \subseteq \mathcal{X}^\omega$, for some space \mathcal{X} . For simplicity, let’s take \mathcal{X} to be \mathbb{N} . The game is played by two players, I and II, who take turns. At each turn, the player chooses a number. This defines an infinite *play* α (an irrational in our case); then player I wins the game if $\alpha \in A$.

A *strategy* for a player is a function from sequences of numbers to numbers: given a play-so-far, it defines a next move. A *winning strategy* for a player is a strategy which, if followed, guarantees winning. A game is *determined* if one player has a winning strategy. It may seem a little surprising, but not all games are determined: if the payoff set A is very complicated, there may not be a winning strategy for either player—unless you add some axioms to ZF . . .

So one can ask, for what sets A is the game on A determined? Write $Det(\Gamma)$ to mean the game on A is determined for all $A \in \Gamma$. Gale and Stewart in their original paper showed that $Det(\Pi_1^0)$, i.e. that all closed games are determined. In the 50s, this was improved to $Det(\Sigma_2^0)$. However, the major advance was made in the early 70s by Martin, who showed $Det(\Delta_1^1)$. He also showed that if MC holds, then $Det(\Sigma_1^1)$.

It turns out that for any significant further advance, one just has to assume the result. There are two such axioms commonly studied: *Projective Determinacy* (PD), which asserts that all the projective classes are determined; and the *Axiom of Determinacy* (AD), which simply asserts that all games are determined! These axioms are very strong: AD contradicts AC (but is consistent with DC), and implies that \aleph_1 is a measurable cardinal; it also implies LM . However, most people would consider it to be false. Despite that, working in the theory $ZF + DC + AD$ can provide interesting and useful results, because even though one might consider AD to be false, one might expect there to be a restricted universe in which it’s true, and then one can get results about sets in that universe.

6 Induction and fixpoints.

Definition by induction is fundamental to all mathematics. The general form of inductive definitions is that you have some operation which, given a collection of objects, gives you a bigger collection of objects, and you want to close under the application of that definition. For example, the set of formulae of first-order arithmetic is formed by starting with the basic formulae, and closing under the operation that forms conjunctions, quantifications, etc. This looks rather like ordinary recursion; and of course, if you want to check that a string is a well-formed formula, you write a recursive function that winds its way down the inductive definition, checking each step.

The inductive definition of well-formed formulae is fairly simple; in particular, it *closes at* ω : once you've applied the construction rule ω times, you don't get any more formulae by applying it again. In general, however, inductive definitions don't necessarily close at ω . Thus one might expect inductive definition to be quite a powerful way of describing sets—and one can certainly ask how powerful, in terms of the notations we already know about. So we'll now consider the power of inductive definition.

Definition 31 Suppose \mathscr{W} is some set, and $\Phi: \wp(\mathscr{W}) \rightarrow \wp(\mathscr{W})$ is a operation taking subsets of \mathscr{W} to subsets of \mathscr{W} , and suppose further that Φ is monotone. The *iterates* of Φ are defined by induction on the ordinals, thus:

$$\Phi^\zeta = \Phi\left(\bigcup_{\xi < \zeta} \Phi^\xi\right)$$

and we write $\Phi^{<\zeta}$ for $\bigcup_{\xi < \zeta} \Phi^\xi$.

Notice that this single definition covers zero (giving $\Phi^0 = \Phi(\emptyset)$), successor (giving $\Phi^{\zeta+1} = \Phi(\Phi^\zeta)$), and limit iterates. In some earlier work on inductive definitions, a different definition of iterate was used, namely to separate out these cases by defining $\underline{\Phi}^0 = \emptyset$, $\underline{\Phi}^{\zeta+1} = \Phi(\underline{\Phi}^\zeta)$ and $\underline{\Phi}^\lambda = \bigcup_{\xi < \lambda} \underline{\Phi}^\xi$ for limit λ . (I'm just adding underlines to distinguish them from our notation.) Although this definition is superficially perhaps more obvious, the notation of the definition is on the whole more attractive. Unfortunately, the underlined notation has been transferred into the modal mu-calculus; but I live in hope of driving it out again.

Exercise 32 (Trivial) Show that the relation between the two notations is: $\underline{\Phi}^\zeta = \Phi^{<\zeta}$, and conversely $\Phi^\zeta = \Phi(\underline{\Phi}^\zeta)$.

It is convenient to use ∞ as a notation for something 'larger than all the ordinals', and write $\Phi^\infty = \Phi^{<\infty} = \bigcup_{\zeta \in \text{Ord}} \Phi^\zeta$. Of course, ∞ doesn't exist, since Ord is a proper class, but this doesn't matter, because as the following result shows, in any given circumstance we can take ∞ to be some sufficiently large ordinal.

Theorem 33 *If Φ is as above, then*

- (1) *if $\zeta < \xi$ then $\Phi^\zeta \subseteq \Phi^\xi$;*
- (2) *for some ordinal κ of cardinality $\leq |\mathscr{W}|$, we have $\Phi^\infty = \Phi^\kappa = \Phi^{<\kappa}$;*
- (3) *Φ^∞ is the least (pre-)fixpoint of Φ .*

Proof. Exercise. (Hints, in case you haven't seen this before: (1) follows straight from monotonicity; (2) follows because otherwise you'd be putting more than $|\mathscr{W}|$ elements into a subset of \mathscr{W} ; and (3) follows from (2) and an easy induction to show that every iterate is contained in

every pre-fixpoint. (A ‘pre-fixpoint’ of Φ is an A such that $\Phi(A) \subseteq A$. It is easy to show that the least pre-fixpoint is also the least fixpoint.)

That defines a form of induction, and allows us to define sets as fixpoints of a monotone operator, but that doesn’t quite fit neatly into our framework of definability, where we would want to define a relation R by $R(x) \Leftrightarrow$ ‘inductive definition given by $\phi(x)$ for some relation ϕ that is built up (by quantification, for example) from already defined relations. To use the inductive definition, we need to produce a Φ from a formula ϕ . Clearly ϕ will have to have a free ‘relation variable’ W so we can feed in the argument of Φ ; but it also has to return a set in some way—and naturally, it can do this by having a free variable w so that it returns the set $\{w \mid \phi(w, \dots, W)\}$. This leads us to the following definition of inductively defined sets.

Definition 34 Suppose $\phi(w, x, W)$ is a relation on $\mathscr{W} \times \mathscr{X} \times \wp(\mathscr{W})$ (so W ranges over subsets of the space \mathscr{W}) that is monotone in W . Define, for each x , a monotone operation Φ_x on $\wp(\mathscr{W})$ by

$$\Phi_x(W) = \{w \mid \phi(w, x, W)\}$$

and define $\phi^\zeta(w, x) \Leftrightarrow w \in \Phi_x^\zeta$. We then say that a relation $R(x)$ is inductively defined by ϕ if there is some w_0 such that $R(x) \Leftrightarrow \phi^\infty(w_0, x)$.

The appearance of w_0 may seem a little strange, but there’s really no way to avoid it in a general framework. (Compare the modal mu-calculus, where fixpoints appear directly without a w_0 ; but modal mu-calculus doesn’t have any first-order variables at all!)

Now we can ask, if ϕ is in some class Γ of relations, what can we say about the relations inductively defined by ϕ ? Let us write $IND(\Gamma)$ for the relations inductively defined by Γ . Of course, we have the catch that ϕ has to be monotone, so Γ had better contain only monotone relations. The easiest way to ensure this is to stick to positive formulae (i.e. W occurs within the scope of an even number of negations)—so if Γ is a class of definable relations, let $pos\text{-}\Gamma$ denote the positive relations in that class.

In order to simplify things, we’ll assume from now on that \mathscr{W} is actually \mathbb{N} , so we’re only taking fixpoints of operations on sets of numbers; after all, taking fixpoints of operations on sets of irrationals feels like a third-order thing, and we certainly don’t want that! (In fact, there are things to say about induction on \mathscr{N} , but we shan’t say them). The following theorems assume $\mathscr{W} = \mathbb{N}$.

In the simplest case, inductive definitions don’t actually buy you anything: induction over a semi-recursive relation achieves nothing.

Theorem 35 $IND(pos\text{-}\Sigma_1^0) = \Sigma_1^0$

Proof. Fixpoints of semi-recursive operations are simple, in the sense that they close at ω . To see this, consider some $w \in \Phi^\omega = \Phi(\Phi^{<\omega})$. Since Φ is semi-recursive, to determine that $w \in \Phi(\Phi^{<\omega})$ requires looking at only a finite number of w_i in $\Phi^{<\omega}$. Each w_i is in some Φ^{n_i} ; but then $w \in \Phi(\Phi^{<1+\max n_i}) = \Phi^{1+\max n_i} \subseteq \Phi^{<\omega}$. So now $w \in \Phi^\infty \Leftrightarrow \exists n. w \in \Phi^n$, and being in Φ^n can be defined by a semi-recursive function, and we’re done.

Exercise 36 (Long) Flesh out this sketch into a rigorous proof.

However, if we do induction over co-semi-recursive relations, things take off very quickly. Kleene's theorem (yes, another one!) on induction states

Theorem 37 $IND(pos-\Pi_1^0) = \Pi_1^1$

Proof. A little trickier than we really want to do here, but not too bad. The \Rightarrow direction is easy, because sets of integers can be coded up as irrationals, and then we just have to say 'for all sets that are fixpoints, w is in the set', which is a Π_1^1 statement. The \Leftarrow direction involves a little more work.

Once we get to Π_1^1 , another induction doesn't get us any further:

Theorem 38 $IND(\Pi_1^1) = \Pi_1^1$

Proof. As noted in the previous theorem, taking least fixpoints is a Π_1^1 operation, so taking least fixpoints over Π_1^1 is still Π_1^1 .

Suppose we negate our least fixpoints, and then do induction? In other words, do induction over a greatest fixpoint; or do nested induction and co-induction. (The greatest fixpoint of Φ is the complement of the least fixpoint of the dual operator $\hat{\Phi}$ defined by $\hat{\Phi}(W) = \neg\Phi(\neg W)$ —prove this, if you haven't already seen it.)

One of the most surprising things to me about the study of inductive definitions is that the obvious hierarchy of nested induction and co-induction was not studied at all until the early 90s. Perhaps a reason for this is that it is actually extremely hard to characterize its power in terms of something else; but as we shall also see, the most basic properties are easy, and not even they were established—or at least, they weren't written down anywhere!

In the classical work on inductive definability, only the following result was stated. Recall that one induction over first-order properties gives us Π_1^1 , which if complemented gives Σ_1^1 .

Theorem 39 $R(w)$ is in $IND(\Sigma_1^1)$ iff R is the set of w such that player I wins the game $\{\alpha \mid P(w, \alpha)\}$, where P is a Σ_2^0 relation.

We shan't prove this, but by the end of this section of the course, we should have a vague woolly feeling as to why it's true.

7 Arithmetic with fixpoints.

7.1 The language and its normal form.

So, since we wish to study mixed induction and coinduction, let us add them to the our logic. Starting with first-order arithmetic as before, we add *set variables* W, X, Y, \dots , new atomic terms of the form $\tau \in \Xi$, where τ is a (first-order) term and Ξ is a *set term*; and set terms are either set variables or have the form $\mu(w, W). \phi$, where ϕ is a formula, w is a first-order variable, and W a set variable; the two variables are bound in the term. We also have the dual set term $\nu(w, W). \phi$, and as usual we'll use duality to push negations inward, and work entirely in positive form. The meaning of the set terms is the obvious one: $\mu(w, W). \phi$ denotes the set $\phi^\infty(w, \dots)$ as previously defined, where \dots denotes the other free variables of ϕ . We'll freely use various obvious notations for the iterates of ϕ , which in modal mu-calculus are traditionally called *approximants* of the fixpoints.

Since we are going to work entirely in positive form, we also need a notation for the ‘iterates’ of maximal fixpoints. These can be defined entirely by duality: $\nu^\zeta(w, W). \phi = \neg \mu^\zeta(w, W). \hat{\phi}$ (where $\hat{\cdot}$ denotes the dual formula), or alternatively can be defined directly by $\nu^\zeta(w, W). \phi = \{ w \mid \phi(w, \bigcap_{\xi < \zeta} \nu^\xi(w, W). \phi) \}$. (Exercise: check these agree!)

As a space saving notational convention, let us write just μX for $\mu(x, X)$, and so on.

It's quite hard to give any non-trivial examples of formulae in this language, so here are a couple of trivial examples. $\mu X. x = 0 \vee (x > 1 \wedge (x - 2) \in X)$ is the set of even numbers; the formula is just writing down the inductive definition of ‘even’. Of course, the even numbers are also the complement of the odd numbers: the odd numbers are defined by $\mu X. x = 1 \vee (x > 1 \wedge (x - 2) \in X)$, so by negating we can express the even numbers as a maximal fixpoint $\nu X. x \neq 1 \wedge (x > 1 \Rightarrow (x - 2) \in X)$. (Here the intuition is that we start with everything, and then throw out the odd numbers one after the other.)

Notice (and verify) that in both these cases, we actually get the same answer whether we take minimal or maximal fixpoints. This is essentially because in both cases we are coding a well-founded inductive definition.

Exercise 40 Consider the formula

$$\mu X. (x = 0 \vee (x + 2) \in X)$$

Show that if μ is μ , then the formula denotes the set $\{0\}$, and if μ is ν , it denotes \mathbb{N} .

This illustrates a general trend: for simple fixpoints, either both are the same, or one of them is trivial. Can you come up with an example formula where the two fixpoints are distinct and non-trivial?

We can now define the obvious hierarchy of fixpoint formulae, along the same lines as before. Of course, we now have set terms in the language, so we have to include them.

Definition 41 If a formula or set term has no fixpoint operators, it is Σ_0^μ and Π_0^μ . Otherwise, for $n \geq 1$:

- If ϕ and Ξ are Σ_{n-1}^μ or Π_{n-1}^μ , they are both Σ_n^μ and Π_n^μ .
- If ϕ_1, ϕ_2 and Ξ are Σ_n^μ , the formulae $\phi_1 \vee \phi_2, \phi_1 \wedge \phi_2, \exists x. \phi_1, \forall x. \phi_1, \tau \in \Xi$ are Σ_n^μ ; and the set term $\mu X. \phi_1$ is Σ_n^μ .
- If ϕ is Σ_n^μ , then the set term $\nu X. \phi$ is Π_{n+1}^μ .

We remarked when dealing with first-order arithmetic that our slightly non-standard definition had the same effect as the usual definition, and the proof of this relied upon the prenex normal form theorem for arithmetic. If the above definition of fixpoint alternation is to behave nicely, there had better be a prenex normal form for arithmetic with fixpoints. Fortunately, there is; but as it is not to be found in standard texts, and is also rather more complicated to prove than the first-order case, we'll give the proof. (Note that the modal mu-calculus does *not* have a prenex normal form, which is why there is some mess involved with the definition of alternation there, as we shall see.) (Most of the rest of this section is lifted from my TCS/CONCUR '96 and STACS '98 papers.)

Definition 42 A Σ_n^μ formula of arithmetic with fixpoints is in *pair-normal form* if it has the form

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

(the last fixpoint being μ or ν according as n is odd or even) where ϕ is first-order.

The terminology 'pair-normal' comes from Lubarsky's paper on mu-arithmetic, and is used because there is a weaker normal form that can be used in structures that can't code up the pairing function. He then proved the normal form theorem. (The presentation here is somewhat better ...)

Theorem 43 If ϕ is Σ_n^μ (Π_n^μ), it is semantically equivalent to a pair-normal formula that is also Σ_n^μ (Π_n^μ).

Proof. We proceed by induction on n , and by structural induction on formulae and set terms.

For a set term $\mu X. \phi$, we assume inductively that ϕ is pair-normal; then we are already pair-normal unless ϕ is $\tau \in \mu Y. \psi$. In that case, the translation pairs up X and Y into W in the natural way, so that $m \in X$ iff $\langle 0, m \rangle \in W$ and $n \in Y$ iff $\langle 1, \langle x, n \rangle \rangle \in W$ (remember that τ, ψ and Y may depend on the individual variable x as well as the set variable X). Note that although Y depends on both x and X , we have only explicitly coded the dependency on x . By standard monotonicity arguments about adjacent fixpoints of the same sign, the dependency on X can be ignored. Thus we translate the original term into

$$\mu W. ((w)_0 = 0 \wedge \langle 1, \langle (w)_1, \tau' \rangle \rangle \in W) \vee ((w)_0 = 1 \wedge \psi')$$

where τ' is obtained from τ by replacing every occurrence of x by $(w)_1$, and ψ' is obtained from ψ by replacing every ' $\rho \in X$ ' by ' $\langle 0, \rho \rangle \in W$ ', and every ' $\rho \in Y$ ' by ' $\langle 1, \langle x, \rho \rangle \rangle \in W$ ', and then every x by $((w)_1)_0$ and every y by $((w)_1)_1$. This procedure clearly preserves the level in the hierarchy. Now, as ϕ was pair-normal, its body ψ was a Π_{n-1}^μ formula; hence the body of $\mu W. \dots$ is Π_{n-1}^μ , and by induction can be transformed into a Π_{n-1}^μ pair-normal formula, and we are done.

Now we consider formulae. For the case $\tau \in \Xi$, inductively transform Ξ to its pair-normal form Ξ' , as in the previous paragraph. Note that if the pairing of adjacent fixpoints above is required, then we need to write $\langle 0, \tau \rangle \in \Xi'$, as τ is supposed to be in X , not W .

The booleans are easy, since $(\tau \in \mu Z. \phi) \wedge \psi$ is equivalent to $\tau \in \mu Z. \phi \wedge \psi$. A little care is needed, though: if we have the conjunction of two fixpoints, one μ and the other ν , we need to put the μ on the outside if we're trying to make it Σ_n^μ , and the ν if we're trying to make it Π_n^μ .

Thus a formula that is both Σ_n^μ and Π_n^μ has a Σ_n^μ pair-normal form and also a Π_n^μ pair-normal form, but does not have a pair-normal form that is both Σ_n^μ and Π_n^μ .

For formulae $\exists x. \phi$, assume that ϕ is $\tau \in \mu Y. \psi$. The existential quantifier is pushed inside the fixpoint by a similar construction to that used in the case of set terms: let W be a new variable, and build ψ' from ψ exactly as before. Then the set term

$$\mu W. (w = \langle 0, 0 \rangle \wedge \exists x. \langle 1, \langle x, \tau \rangle \rangle \in W) \vee ((w)_0 = 1 \wedge \psi')$$

contains $\langle 0, 0 \rangle$ iff $\exists x. \phi$. Now the case of ϕ being $\tau \in \nu Y. \psi$ is similar.

Similarly for formulae $\forall x. \phi$.

So we see that the transformation makes no change to the Σ_n^μ level, as claimed. \square

Exercise 44 (Long) Fill in the details; in particular, prove the correctness of the ‘push \exists through μ ’ construction.

This is quite a complex construction, and some simple examples may be helpful. Firstly, consider ‘ t is even and t is not a multiple of three’. If we use an inductive definition of ‘multiple of three’ (quite unnecessary, of course, but never mind), we get

$$(t \in \mu X. x = 0 \vee (x > 1 \wedge (x - 2) \in X))$$

$$\wedge (t \in \nu Y. y \neq 0 \wedge (y > 2 \Rightarrow y - 3 \in Y)).$$

As the two fixpoints are independent, we can move one inside the other to get

$$t \in \mu X. t \in \nu Y. (x = 0 \vee (x > 1 \wedge (x - 2) \in X))$$

$$\wedge (y \neq 0 \wedge (y > 2 \Rightarrow y - 3 \in Y)).$$

As an example of the treatment of first-order quantifiers, consider ‘ t is a composite number’. Let us again, for the purposes of exposition, use an inductive definition of multiple, but use an existential quantifier over possible factors, that is to say ‘there is an $x > 1$ such that t is a multiple (> 1) of x ’:

$$\exists x. x > 1 \wedge t \in \mu Y. y = 2x \vee (y > 2x \wedge (y - x \in Y)).$$

Applying the construction given above yields, where for readability we write w_{10} etc. for $((w)_1)_0$ etc.:

$$\langle 0, 0 \rangle \in \mu W. (w = \langle 0, 0 \rangle \wedge \exists x. \langle 1, \langle x, t \rangle \rangle \in W)$$

$$\vee (w_0 = 1 \wedge w_{10} > 1 \wedge (w_{11} = 2w_{10} \vee$$

$$(w_{11} > 2w_{10} \wedge \langle 1, \langle w_{10}, w_{11} - w_{10} \rangle \rangle \in W))).$$

Here the meat of the inductive definition is the same as before, but it’s now being carried on in the $((\)_1)_1$ component of W , which is parametrized by the $((\)_1)_0$ component representing x . The first line says, effectively, that the flag value $\langle 0, 0 \rangle$ is in W only if $\exists x. t \in \mu Y. \dots$, and the second and third lines compute Y as the last component of W , with the constraint on x included in this computation.

Exercise 45 It was claimed in the lecture that existential quantification can be replaced by a least fixpoint, though I didn’t write it down correctly. It’s true, though: prove that

$$(\exists x. \phi) \Leftrightarrow 0 \in \mu X. \phi \vee (x + 1) \in X$$

(Of course, this doesn't help in general, since we don't want to replace simple first-order quantifiers by fixpoints, but it's a simpler illustration of the 'flag value' idea.)

7.2 The hierarchy theorem.

Having established the normal form theorem, we can now deal with the hierarchy. It would in fact be possible to proceed exactly as in the previous hierarchy theorems, taking the Kleene theorems as the base case; but there are one or two minor complications to worry about. It turns out to be easier to prove the entire theorem from scratch, via a rather pleasant encoding.

The strategy is the same as before, namely to show that the truth of Σ_n^μ formulae can itself be expressed by a Σ_n^μ formula, and to use a diagonalization argument to show that this formula cannot be equivalent to any Π_n^μ formula.

Firstly, take a suitable Gödel numbering of mu-arithmetic. We consider only formulae without free set variables; wlog, we may assume that all encoded formulae are in normal form, and are normalized so that the free individual variables are z_0, \dots, z_k , the first-order quantifiers bind z_{k+1}, \dots , and for a formula of alternation depth n , the fixpoint variables are X_n, \dots, X_1 , with associated individual variables x_n, \dots, x_1 . We use sans-serif type to indicate that the variable is being seen as part of an encoded object-level formula; normal italic type indicates a meta-level variable. We use corner quotes to denote the Gödel coding. We also need coded *assignments* which map an encoded variable to a value: we write $[v/z]$ for the assignment that maps z (strictly, the code $\ulcorner z \urcorner$) to the integer v , and $a[v/z]$ for the updating of a by $[v/z]$. We use double quotes to indicate the appropriate meta-language formalization of the informal statement inside the quotes.

Now suppose that $\text{Sat}_n(x, y)$ is a formula of mu-arithmetic expressing the truth of Σ_n^μ formulae, so that if ϕ is a formula and a an assignment of values \vec{v} to the free variables \vec{z} of ϕ , then $\text{Sat}_n(\ulcorner \phi \urcorner, a)$ is true just in case $\phi(\vec{v}/\vec{z})$ is true. We have the

Lemma 46 $\text{Sat}_n(z_0, [z_0/z_0])$ is not equivalent to any Π_n^μ formula.

Proof. The proof is exactly as for the arithmetical hierarchy. Suppose the contrary, i.e. that $\neg \text{Sat}_n(z_0, [z_0/z_0])$ is equivalent to some Σ_n^μ formula $\theta(z_0)$. Then we have

$$\theta(\ulcorner \theta \urcorner) \text{ iff } \neg \text{Sat}_n(\ulcorner \theta \urcorner, [\ulcorner \theta \urcorner/z_0]) \text{ iff } \neg \theta(\ulcorner \theta \urcorner)$$

, which is a contradiction. □

It remains to show that Sat_n exists and is indeed a Σ_n^μ formula.

Theorem 47 Sat_n is a Σ_n^μ formula of mu-arithmetic, for $n > 0$.

Proof. We start by constructing Sat_0 , truth in first-order arithmetic, both as a Σ_1^μ formula and as a Π_1^μ formula. $\text{Sat}_0(x, y)$ is defined as:

$$\begin{aligned} \langle x, y \rangle \in \mu(w, W). \quad & \text{“}(w)_0 = \ulcorner P(\tau) \urcorner \text{ and } \text{pred}(\ulcorner P \urcorner, \text{eval}(\ulcorner \tau \urcorner, (w)_1))\text{”} \\ & \vee \text{“}(w)_0 = \ulcorner \phi_1 \wedge \phi_2 \urcorner \text{ and } (\langle \ulcorner \phi_1 \urcorner, (w)_1 \rangle \in W \wedge \langle \ulcorner \phi_2 \urcorner, (w)_1 \rangle \in W)\text{”} \\ & \vee \text{“}(w)_0 = \ulcorner \phi_1 \vee \phi_2 \urcorner \text{ and } (\langle \ulcorner \phi_1 \urcorner, (w)_1 \rangle \in W \vee \langle \ulcorner \phi_2 \urcorner, (w)_1 \rangle \in W)\text{”} \\ & \vee \text{“}(w)_0 = \ulcorner \exists z_i. \phi_1 \urcorner \text{ and } \exists v. \langle \ulcorner \phi_1 \urcorner, (w)_1[v/z_i] \rangle \in W\text{”} \\ & \vee \text{“}(w)_0 = \ulcorner \forall z_i. \phi_1 \urcorner \text{ and } \forall v. \langle \ulcorner \phi_1 \urcorner, (w)_1[v/z_i] \rangle \in W\text{”} \end{aligned}$$

where $\text{eval}(t, a)$ is the recursive function which evaluates a coded term $t = \ulcorner \tau \urcorner$ under the variable assignment a , and $\text{pred}(p, x)$ is the computable predicate which is true if the value x satisfies the predicate coded by $p = \ulcorner P \urcorner$.

We have here skipped the details of the coding, which are standard. For example, if we look in more detail at the clause for \forall , it actually says:

$$f((w)_0) = \ulcorner \forall v. \langle g((w)_0), h((w)_1, v, g'((w)_0)) \rangle \in W$$

where f extracts the top-level connective of a coded formula, g extracts the body of a \forall formula and g' extracts the bound variable, and $h(a, v, z)$ takes the variable assignment a and updates the variable whose code is z by the value v . The fact that these functions f, g, h are recursive is obvious, and since we allow ourselves all recursive functions as primitives, that is sufficient; but explicit definitions in standard arithmetic may be found in standard references.

It is clear that this fixpoint formula simply encodes directly the definition of truth in arithmetic. The formula is Σ_1^μ , but since the encoded recursive function terminates on all arguments—it is just a definition by induction on the structure of formulae—it does not matter whether we use a minimal or maximal fixpoint to achieve the recursion. Thus we may also obtain Sat_0 as a Π_1^μ formula.

In order to encode within mu-arithmetic the evaluation of formulae with fixpoints, it is necessary to have the same fixpoint structure in the Sat formula as in the formula it's evaluating. Recall that we assume pair-normal form, and suppose that we wish to evaluate Σ_n^μ formulae where n is odd, that is, formulae of the form

$$\tau_n \in \mu X_n. \tau_{n-1} \in \nu X_{n-1} \dots \tau_2 \in \nu X_2. \tau_1 \in \mu X_1. \phi \quad (*)$$

where ϕ is first-order. The interpretation of the pure first-order part of ϕ may be done with the Σ_1^μ version of Sat_0 —but ϕ may also now contain formulae $\tau \in X_i$. We cannot code as integers the sets referred to by the X_i , so they must be represented by set variables in the meta-language. We use the meta-level variable X_i to represent the object variable X_i , and extend the body of Sat_0 by the clauses (for each $1 \leq i \leq n$)

$$\vee \text{“}(w)_0 = \ulcorner \tau \in X_i \urcorner \text{ and } \text{eval}(\ulcorner \tau \urcorner, (w)_1) \in X_i \text{”}.$$

Let Sat'_0 denote the adjusted Sat_0 .

With these adjustments, we have that $(*)$ is true with free variable assignment a just in case

$$\begin{aligned} \text{eval}(\ulcorner \tau_n \urcorner, a) &\in \mu X_n. \\ \text{eval}(\ulcorner \tau_{n-1} \urcorner, a[x_n/x_n]) &\in \nu X_{n-1}. \dots \\ \text{eval}(\ulcorner \tau_1 \urcorner, a[x_n, \dots, x_2/x_n, \dots, x_2]) &\in \mu X_1. \\ \text{Sat}'_0(\ulcorner \phi \urcorner, a[x_n, \dots, x_1/x_n, \dots, x_1]) & \end{aligned}$$

Now we just parametrize on $(*)$: let $f_1(x, y)$ be the function that given x encoding a Σ_n^μ formula $(*)$ and an assignment y , computes $\text{eval}(\ulcorner \tau_n \urcorner, y)$, and so on, and let $g(x)$ extract the body of $(*)$. Then we have $\text{Sat}_n(x, y)$ in the form

$$\begin{aligned} f_n(x, y) &\in \mu X_n. f_{n-1}(x, y[x_n/x_n]) \in \nu X_{n-1}. \dots \\ f_1(x, y, [x_n, \dots, x_2/x_n, \dots, x_2]) &\in \mu X_1. \text{Sat}'_0(g(x), y[x_n, \dots, x_1/x_n, \dots, x_1]) \end{aligned}$$

which is Σ_n^μ as required. If n is even, we use the Π_1^μ version of Sat_0 instead.

The fact that Sat_n does indeed code truth is easily shown: show by induction on i that each meta-level fixpoint set X_i coincides with the object-level set X_i . The base case follows from the correctness of Sat'_0 , and the induction step is easy.

It may be noted that we have also skipped details of what the functions f_1 etc. should do if given ill-formed arguments. Any convenient trick may be used; the details are unimportant. \square

That concludes our investigation of arithmetic with fixpoints; the next stage is the somewhat surprising discovery that it can be used to solve a long-standing problem in the modal mu-calculus.

8 Alternation in the modal mu-calculus.

8.1 The language; simple alternation.

I assume that you are familiar with the modal mu-calculus, so let us just briefly review the syntax and semantics.

We assume some countable set \mathcal{L} of *labels*. The formulae Φ are defined inductively thus: variables Z and the boolean constants tt , ff are formulae; if Φ_1 and Φ_2 are formulae, so are $\Phi_1 \vee \Phi_2$ and $\Phi_1 \wedge \Phi_2$; if Φ is a formula and l a label, then $[l]\Phi$ and $\langle l \rangle \Phi$ are formulae; and if Φ is a formula and Z a variable, then $\mu Z. \Phi$ and $\nu Z. \Phi$ are formulae. (As with first-order logic, we adopt the convention that the scope of the binding operators μ and ν extends as far as possible. Note also that in this section we'll stick to the convention that upper-case Greek letters Φ, Ψ, Υ are modal mu-calculus formulae, whereas lower-case Greek letters ϕ, ψ are mu-arithmetic formulae.)

Given a labelled transition system $\mathcal{T} = (\mathcal{S}, \mathcal{L}, \longrightarrow)$, where \mathcal{S} is a set of states, \mathcal{L} a set of labels, and $\longrightarrow \subseteq \mathcal{S} \times \mathcal{L} \times \mathcal{S}$ is the transition relation (we write $s \xrightarrow{l} s'$), and given also a valuation \mathcal{V} assigning subsets of \mathcal{S} to variables, the denotation $\|\Phi\|_{\mathcal{V}}^{\mathcal{T}} \subseteq \mathcal{S}$ of a mu-calculus formula Φ is defined in the obvious way for the variables and booleans, for the modalities by

$$\begin{aligned} \|[l]\Phi\|_{\mathcal{V}}^{\mathcal{T}} &= \{ s \mid \forall s'. s \xrightarrow{l} s' \Rightarrow s' \in \|\Phi\|_{\mathcal{V}}^{\mathcal{T}} \} \\ \|\langle l \rangle \Phi\|_{\mathcal{V}}^{\mathcal{T}} &= \{ s \mid \exists s'. s \xrightarrow{l} s' \wedge s' \in \|\Phi\|_{\mathcal{V}}^{\mathcal{T}} \}, \end{aligned}$$

and for the fixpoints by

$$\begin{aligned} \|\mu Z. \Phi\|_{\mathcal{V}}^{\mathcal{T}} &= \bigcap \{ S \subseteq \mathcal{S} \mid \|\Phi\|_{\mathcal{V}[Z:=S]}^{\mathcal{T}} \subseteq S \} \\ \|\nu Z. \Phi\|_{\mathcal{V}}^{\mathcal{T}} &= \bigcup \{ S \subseteq \mathcal{S} \mid S \subseteq \|\Phi\|_{\mathcal{V}[Z:=S]}^{\mathcal{T}} \}. \end{aligned}$$

Approximants (iterates) of the fixpoint formulae are defined just as before.

Note a major difference from mu-arithmetic: there is no distinction between formulae and set terms, because all formulae are implicitly describing sets of states.

To start with, we'll define the hierarchy of alternating formulae by exact analogy with arithmetic. We'll see later that this is not really what we want to do; but that will be a detail we can resolve independently of the main argument.

Definition 48 A formula with no fixpoint operators is $\Sigma_0^{S\mu}$ and $\Pi_0^{S\mu}$. Otherwise, for $n \geq 1$:

- If Φ is $\Sigma_{n-1}^{S\mu}$ or $\Pi_{n-1}^{S\mu}$, it is both $\Sigma_n^{S\mu}$ and $\Pi_n^{S\mu}$.
- If Φ_1 and Φ_2 are $\Sigma_n^{S\mu}$, so are $\Phi_1 \vee \Phi_2$, $\Phi_1 \wedge \Phi_2$, $\langle l \rangle \Phi_1$, $[l]\Phi_1$, and $\mu X. \Phi_1$.
- If Φ is $\Sigma_n^{S\mu}$, then $\nu X. \Phi$ is $\Pi_{n+1}^{S\mu}$.

The S here stands for 'simple' or 'syntactic'; later, we'll define stronger notions of alternation.

In the model-checking work on alternation, you won't see alternation expressed in terms of these classes of formulae. Instead, you'll see references to the 'alternation depth' of a formula. This is defined algorithmically, but we can cast it in terms of our classes:

Definition 49 The *simple alternation depth* $\text{ad}^S(\Phi)$ of a formula Φ is the least n such that Φ is both $\Sigma_n^{S\mu}$ and $\Pi_n^{S\mu}$.

The motivation for this definition gives one of the main reasons why it has been considered interesting to study the alternation hierarchy.

Theorem 50 *If Φ is a formula and \mathcal{T} a transition system, with $|\mathcal{T}| \cdot |\Phi| = n$ and $\text{ad}(\Phi) = d$, then all known algorithms for model-checking, i.e. determining whether $s \in \|\Phi\|$, have worst-case complexity $\Omega(n^{\Omega(d)})$.*

(This is actually true for the stronger notion of alternation depth that we'll see later, not for the simple one; hence I've dropped the S .) In other words, all known algorithms for model-checking are at least exponential in the alternation depth. However, although the best algorithms also give us an $O(n^{O(d)})$ upper bound for the complexity of the model-checking problem, we don't have any useful lower bounds. (The problem is known to be PTIME-hard, but that's no great surprise.) In particular, we don't know whether the problem is polynomial or not. We do know that it is both NP and co-NP, which makes it one of very few examples of such problems that are not known to be polynomial (primality is another). It seems possible that the above is also a lower bound; but that would imply $P \neq NP$, so is unlikely to be easy to prove. Of course, if the alternation hierarchy collapsed, one might hope to prove that the problem is polynomial, by smashing all formulae down to low alternation equivalents. However, the alternation hierarchy doesn't collapse.

8.2 The strictness of the simple alternation hierarchy.

The strategy here is to transfer the result from mu-arithmetic. Specifically, we shall show that on 'reasonable' transition systems, the denotation of a $\Sigma_n^{S\mu}$ formula is (coded as) a Σ_n^μ set of integers; and conversely, we shall construct a reasonable transition system and a $\Sigma_n^{S\mu}$ formula Φ such that the denotation of Φ is a strict Σ_n^μ set of integers, and so Φ can't be equivalent to any lower alternation formula.

We define a *recursively presented transition system (r.p.t.s.)* to be a labelled transition system $(\mathcal{S}, \mathcal{L}, \longrightarrow)$ such that \mathcal{S} is (recursively codable as) a recursive set of integers, \mathcal{L} likewise, and \longrightarrow is recursive. Henceforth we consider only recursively presented transition systems, with recursive valuations for the free variables. We have the following theorem:

Theorem 51 *For a modal mu-calculus formula $\Phi \in \Sigma_n^{S\mu}$, the denotation $\|\Phi\|$ in any r.p.t.s. is a $\Sigma_n^{S\mu}$ definable set of integers.*

Proof. This is a trivial translation of the semantics of the modal mu-calculus into arithmetic. For each modal formula Ψ , we define an arithmetic formula $\psi(s)$ such that $\psi(s) \Leftrightarrow s \in \|\Psi\|$, by structural induction. For example, the translation of $\mu X. \Psi$ is $s \in \mu(x, X). \psi(x)$. \square

The remaining task is to construct an r.p.t.s. and a modal formula with a strict Σ_n^μ denotation. The trick here is to build a machine which 'interprets' formulae of mu-arithmetic. Now, a recursively presented transition system can't possibly correctly interpret incredibly complicated things like mu-arithmetic, so we'll cheat, and move part of the interpretation into a modal formula, which tells us which execution sequences of the machine are 'correct' interpretations. Roughly, the easy recursive parts of the job are done by the transition system, the first-order quantifiers are specified by modalities, and the arithmetic fixpoints are specified by modal fixpoints. So, let's proceed.

We aim to construct a transition system \mathcal{T} and a $\Sigma_n^{S\mu}$ modal mu-calculus formula Φ such that the set of states satisfying Φ is defined by the strict Σ_n^μ arithmetic formula Sat_n .

The transition system \mathcal{T} should be viewed as a machine for evaluating arithmetic expressions in the same way that Sat_n does: the computation happening in the body of Sat'_0 will be dealt with by the definition of the transitions of \mathcal{T} , and the arithmetic fixpoints are translated into modal fixpoints in Φ .

The states of \mathcal{T} encode several pieces of information. Namely, a state s contains: a formula ψ_s of the form $(*)$ (see page 23), and a variable assignment a_s , and a pointer p_s into ψ_s which keeps track of where we are in the evaluation. We use the notation of $(*)$ to refer to parts of ψ_s .

The labels of \mathcal{T} are used to distinguish various steps of computation; we shall start with enough labels to make the construction clear, and then argue the number down a little.

The transitions of \mathcal{T} from a state s are defined thus:

- If p_s points at τ_i (or after μX_{i+1} , which we consider to be the same), then $s \xrightarrow{x_i} s'$ where $\psi_{s'} = \psi_s$, and $a_{s'} = a_s[\text{eval}(\tau_i, a_s)/x_i]$ and $p_{s'}$ points after μX_i . That is, the term τ_i is evaluated in the current assignment, x_i is set to its value, and we start evaluating the inner fixpoint.

Otherwise, p_s points at a subformula of ϕ . The transition from s mimics the appropriate clause of Sat'_0 . The ψ component is not altered by any transition, and the a component is unchanged unless otherwise stated.

- If p_s points at $P(\tau)$, then $s \xrightarrow{a} s_a$ ('a' for atom), where s_a is a special state with no structure, only if $P(\tau)$ is true with variable assignment a_s ; otherwise there are no transitions from s .
- If p_s points at $\phi_1 \wedge \phi_2$, then $s \xrightarrow{c} s_k$ ('c' for conjunction) for $k = 1, 2$, where p_{s_k} points at ϕ_k .
- If p_s points at $\forall z_i. \phi_1$, then $s \xrightarrow{c} s_k$ (universal quantification is treated as conjunction) for $k \in \mathbb{N}$, where p_{s_k} points at ϕ_1 , and $a_{s_k} = a_s[k/z_i]$.
- If p_s points at $\phi_1 \vee \phi_2$, then $s \xrightarrow{d} s_k$ ('d' for disjunction) for $k = 1, 2$, where p_{s_k} points at ϕ_k .
- If p_s points at $\exists z_i. \phi_1$, then $s \xrightarrow{d} s_k$ (existential quantification is treated as disjunction) for $k \in \mathbb{N}$, where p_{s_k} points at ϕ_1 , and $a_{s_k} = a_s[k/z_i]$.
- If p_s points at $\tau \in X_i$, then $s \xrightarrow{x_i} s'$, where $p_{s'}$ points after μX_i , and $a_{s'} = a_s[\text{eval}(\tau, a_s)/x_i]$. That is, the term τ is evaluated, copied to the input variable x_i of the fixpoint X_i , and evaluation of the fixpoint started.

It is clear that \mathcal{T} is a recursively presented transition system.

Now consider the following *modal* mu-calculus formula:

$$\begin{aligned} \text{MuSat}_n &\stackrel{\text{def}}{=} \langle x_n \rangle \mu X_n. \langle x_{n-1} \rangle \nu X_{n-1}. \dots \langle x_1 \rangle \mu X_1. \mu W. \\ &\quad \langle a \rangle \text{tt} \vee (\langle c \rangle \text{tt} \wedge [c]W) \vee \langle d \rangle W \\ &\quad \vee \langle x_1 \rangle X_1 \vee \dots \vee \langle x_n \rangle X_n \end{aligned}$$

By the construction of \mathcal{T} , we have:

Theorem 52 $s \models \text{MuSat}_n$ just in case p_s points at ψ_s , and $\text{Sat}_n(\ulcorner \psi_s \urcorner, a_s)$. Hence MuSat_n is a strict $\Sigma_n^{S\mu}$ modal formula.

Proof. A fairly routine inductive proof. It's a simplified version of the main theorem in my TCS paper on the topic. \square

MuSat_n is already quite a simple formula, but it is interesting to try to simplify it further, which we shall do in stages.

Firstly, is it necessary to have the double occurrence of $\langle x_i \rangle$, or can we remove the guards from the fixpoint formulae? Yes, we can: consider the formula

$$\text{MuSat}'_n \stackrel{\text{def}}{=} \mu X'_n. \nu X'_{n-1}. \dots \mu X'_1. \mu W. \Psi$$

where Ψ is formed from the body of MuSat_n by priming the X_i s. The relation between MuSat_n and MuSat'_n is that $X'_n = \dots = X'_1 = X_1$ (note that in MuSat_n, we have $X_1 \supseteq X_2 \cup \dots \cup X_n$), and conversely $X_i = \langle x_i \rangle X'_i$ for $i = n, \dots, 2$. The denotation of MuSat'_n is still a strict Σ_n^μ set, since the denotation of MuSat_n is a projection of it.

Next, the occurrence of $\langle c \rangle \text{tt}$ is irritating. Its purpose is to assert that p_s is indeed pointing at an \wedge -subterm of ψ_s —of course, $[c]W$ is true at any state with no c -transitions from it. However, we can render it unnecessary by modifying \mathcal{T} : if s is any state *other than* an \wedge -subterm state, then add a transition $s \xrightarrow{c} s$. Since W is a least fixpoint variable, if W is true at a state with a c -loop, it is true by virtue of some other disjunct than $[c]W$, and it is not true if it was not true before the loop was added.

We can also eliminate the requirement for a separate a -transition, by modifying the modification: for all those states s with an a -transition, remove the c -loop added in the previous paragraph; now $[c]W$ is true at those states, so we can discard the $\langle a \rangle \text{tt}$ clause.

Finally, we note that $W = X'_1$, and they are adjacent least fixpoints, so we can amalgamate them; further, the job of the d transition can as well be done by x_1 , since they work on disjoint sets of states.

Hence we arrive at the following very simple example of a strict $\Sigma_n^{S\mu}$ modal formula (replacing X' by X again):

$$\text{MuSat}''_n \stackrel{\text{def}}{=} \mu X_n. \nu X_{n-1}. \dots \mu X_1. [c]X_1 \vee \langle x_1 \rangle X_1 \vee \dots \vee \langle x_n \rangle X_n$$

8.3 ‘Real’ alternation.

Finally, let us return to the question of alternation in the modal mu-calculus, and how it should be defined. The reason for questioning the utility of the simple definition is this: consider the formula $\mu X. (\nu Y. P \wedge [a]Y) \vee \langle b \rangle X$. This is $\Sigma_2^{S\mu}$, and not $\Pi_2^{S\mu}$, and so its simple alternation depth is 2; but in terms of complexity, it’s no harder to compute than a formula with just one fixpoint, because the inner fixpoint is entirely independent of the outer fixpoint: so it ‘should’ have alternation depth 1. On the other hand, $\mu X. \nu Y. (P \wedge [a]Y) \vee \langle b \rangle X$ does have an interdependence, and is hard to compute. Hence when the notion of alternation depth was defined by the model-checking people, it was done like this:

Definition 53 The *Emerson–Lei alternation classes and depth* $\Sigma_n^{EL\mu}$ and ad^{EL} are defined as in definition 48, with the addition of the following clause:

- if $\Phi(X)$ is $\Sigma_n^{EL\mu}$ with free variable X , and Ψ is $\Sigma_n^{EL\mu}$, then $\Phi[\Psi/X]$ is $\Sigma_n^{EL\mu}$ *provided that* Ψ *is a closed formula.*

This captures the idea that completely independent fixpoints shouldn’t count as alternation. If you read papers on model-checking, when they talk about alternation depth, they usually mean ad^{EL} . However, this notion is not really satisfactory. Here’s one reason: consider

$$\Upsilon = \mu X_1. \nu X_2. \mu X_3. \dots \nu X_{n-1}. \mu X_n. X_1 \vee X_n.$$

This formula is strict $\Sigma_n^{EL\mu}$, but it's obviously not really complicated, since in fact the middle fixpoints are irrelevant. Another reason is that it is often convenient to consider *simultaneous fixpoints*, where one takes a vector of variables and a vector of defining equations, and takes fixpoints simultaneously over all variables—and the Emerson–Lei definition is not robust when one moves to such fixpoints.

A definition that addresses these problems is that of Niwiński. It works thus:

Definition 54 The *Niwiński alternation classes and depth* $\Sigma_n^{N\mu}$ and ad^N are defined as in definition 48, with the addition of the following clause:

- if $\Phi(X)$ is $\Sigma_n^{N\mu}$ with free variable X , and Ψ is $\Sigma_n^{N\mu}$, then $\Phi[\Psi/X]$ is $\Sigma_n^{EL\mu}$ *provided that no free variable of Ψ is captured by a fixpoint operator of Φ .*

Exercise 55 Show that $\text{ad}^N(\Upsilon) = 1$.

Now it turns out that we can establish the hierarchy for the Niwiński classes as well. Obviously any $\Sigma_n^{S\mu}$ formula is also $\Sigma_n^{N\mu}$, so the hard formulae stay the same. To get the theorem, we need the following:

Theorem 56 *On an r.p.t.s., any $\Sigma_n^{N\mu}$ formula has a denotation that is arithmetic Σ_n^μ .*

Proof. The proof works by showing that the translation into arithmetic produces a formula that can be put into Σ_n^μ . The details involve manipulation along the lines of the normal form theorem, and can be found in the TCS paper. \square

That concludes our investigation of the modal mu-calculus alternation hierarchy.