Ordinal Analysis by Transformations

Henry Towsner

Department of Mathematics, Carnegie Mellon University

Abstract

The technique of using infinitary rules in an ordinal analysis has been one of the most productive developments in ordinal analysis. Unfortunately, one of the most advanced variants, the Buchholz Ω_μ-rule, does not apply to systems much stronger than Π^1_1-comprehension. In this paper, we propose a new extension of the Ω rule using game-theoretic quantifiers. We apply this to a system of inductive definitions with the strength of a recursively inaccessible ordinal.

Key words: Ordinal analysis, inductive definitions
1991 MSC: 03F05, 03F35

1 Introduction

Infinitary inference rules have been a key tool in ordinal analysis since their introduction by Schütte [1]. The appropriate infinitary rule for Peano Arithmetic, the ω rule, is reasonably straightforward—it simply branches over the natural numbers—but suitable infinitary rules for stronger systems are less clear.

The first type proposed, Buchholz's Ω_μ rule [2], branches not over numbers, but over a particular class of derivations. Subsequently, Pohlers proposed the method of local predicativity [3], in which infinitary rules branch over infinite ordinals. Rules branching over ordinals have almost entirely replaced the Ω_μ rule, in large part because they led to productive generalizations, culminating in an analysis of Π^1_2-comprehension [4], while the Ω_μ rule seemed limited to iterated systems of Π^1_1-comprehension.

In the method of local predicativity, ordinals are built directly into the system, since they are necessary to even describe the system cut-elimination will take

Email address: henryt@cmu.edu (Henry Towsner).

place in. This integration with ordinals is different from earlier analyses, in which the cut-elimination process came first and the ordinals could be "read off" from the reduction procedure; in local predicativity, the crucial collapsing step is justified by reference to the properties of the ordinals, which have, naturally, been defined just so as to make this possible. Unfortunately, as the systems get more complex, this leads to the appearance that the proof proceeds by "magic", obscuring the underlying structure of the argument. This problem isn't intrinsic to infinitary techniques—the most advanced finitary methods, as in [5],[6], and [7], also require systems defined in terms of ordinals, and face the same problems as a result.

In the author's opinion, reductions which can be defined independently of ordinals are clearer and have a greater potential for extracting combinatorial consequences. Unfortunately, ordinal based methods have been the only option for going beyond iterations of Π^1_1-comprehension. In this paper, we propose an alternate method of analyzing strong subsystems of analysis, based on a "game-theoretic" extension of the Ω_μ rule, and apply it to a lightface version of the system μ_2 described by Mollerfeld [12]. The exact strength of this system, as characterized by subsystems of analysis or recursively large ordinals, is not known to the author, but Mollerfeld's work implies that it is at least as strong as the system with a recursively inaccessible ordinal, although it may be stronger. Systems with the strength of a recursively inaccessible ordinal were first analyzed using finitary methods in the form of $(\Delta^1_2-CA)_0+BI$ [8] and later using infinitary methods in the form of KPi [9]. For the equivalence between these systems, see [10]. (There is also an analysis using Ω_μ rules, [11], but this uses local predicativity-style ordinal indices to obtain sufficiently large iterations of the Ω_μ rule.)

Systems of inductive definitions are relatively susceptible to ordinal analysis, so we will extend the particularly elegant analysis of $ID_{<\omega}$ given in [13]. We will work with the simplest fixed point operator which can't be analyzed by that method, namely a fixed point of the form

$$\mu xX.A(x,X,\mu yY.B(x,X,y,Y))$$

where Z appears negatively in $A(x,X,Z)$.

Such definitions contain objects of a corecursive character, so it is not surprising that we use the method of corecursion (as described in [14] and [15]) in a key definition.

While this illustrates the method, it by no means exhausts it. Unsurprisingly, since any proof has to exceed methods available in Π^1_1-comprehension—which includes recursion along any easily definable well-ordering—it becomes difficult to even describe the iterated form of the method required to analyze stronger fixed point. Even with this limitation, it appears that, at a minimum, this method extends to the μ_2 calculus (in which the A above could contain closed fixed points of the same form, which could themselves contain fixed points of that form, and so on), and we hope that a suitable generalization would extend to the complete μ-calculus, which is known to be equivalent to Π^1_2-comprehension [12].

2 Outline

Before launching into the technical details of the proof, we outline the derivation of the general method from Buchholz's Ω_μ rule. Suppose that we can prove cut-elimination for some arbitrary theory T (say, Peano Arithmetic or ID_n) using an infinitary system T^∞. We may extend T to a theory T' by adding a least fixed point predicate

$$\mu x X.A(x, X)$$

where A is a formula of T and X appears positively, along with closure and induction axioms. We may then extend T^∞ by a closure rule

$$\frac{A(n, \mu x X.A(x, X))}{n \in \mu x X.A(x, X)}$$

and call proofs in this extended system "small". The full infinitary version of T' adds an Ω rule which branches over small proofs of $n \in \mu x X.A(x, X)$ and gives the conclusion $n \notin \mu x X.A(x, X)$. Cut-elimination is quite easy to prove, and the heart of the resulting argument is the demonstration that the induction axiom in the finitary system can be embedded as an Ω rule in the infinitary system. The proof that this is possible involves showing that, given any "small" proof of $n \in \mu x X.A(x, X)$, the predicate $y \in \mu x X.A(x, X)$ can be systematically replaced by any formula $F[y]$.

This method breaks down when we attempt to add the predicate

$$\mu x X.A(x, X, \mu y Y.B(x, X, y, Y))$$

where X appears negatively in B, and therefore $\mu y Y.B(x, X, y, Y)$ must appear negatively in A. (For convenience, we abbreviate $\mu y Y.B(x, X, y, Y)$ by $\mu_B(x, X)$ and $\mu x X.A(x, X, \mu_B(x, X))$ by μ_A.) If we attempt the same technique, a "small" proof of $n \in \mu_A$ must contain negative occurrences of $\mu_B(m, \mu_A)$, which must be introduced by an Ω rule for $\mu_B(m, \mu_A)$, which must in turn branch over proofs containing negative occurrences of μ_A, which gives a vicious cycle.

To find a way out of this dilemma, we can consider what we expect to happen when we attempt to embed an induction axiom for μ_A as a hypothetical Ω rule.

We would expect to replace μ_A with some formula F, and therefore whatever rule introduces $\neg\mu_B(m, \mu_A)$ must be easily converted to a proof of $\neg\mu_B(m, F)$. This is not true of the ordinary Ω rule, which would face the obstacle that the Ω rule for $\neg\mu_B(m, \mu_A)$ does not even necessarily branch over the right domain to become an Ω rule for $\neg\mu_B(m, F)$.

We resolve both these problems at once by introducing a new type of Ω rule to be the "small" rule introducing $\neg\mu_B(m, \mu_A)$; this rule will branch over proofs of $\mu_B(m, F)$ for any F. This difficult is that such derivations may contain inference rules which more widely than is permitted in a small proof (for instance, the introduction of $\neg\mu_A$ when F is μ_A).

Such inferences will be converted into non-branching inference rules. We will call these inference rules *truncated inferences*, since rather than encoding the manner in which the original proof derived $\neg F[n]$, they merely note where such a derivation occured. The Ω rule will then provide, to each derivation of $n \in \mu_B(m, F)$, a derivation of some $G[n, F]$ from instances of these truncated inferences, as well as an indication, for each truncated inference in the resulting derivation, a source inference in the original derivation.

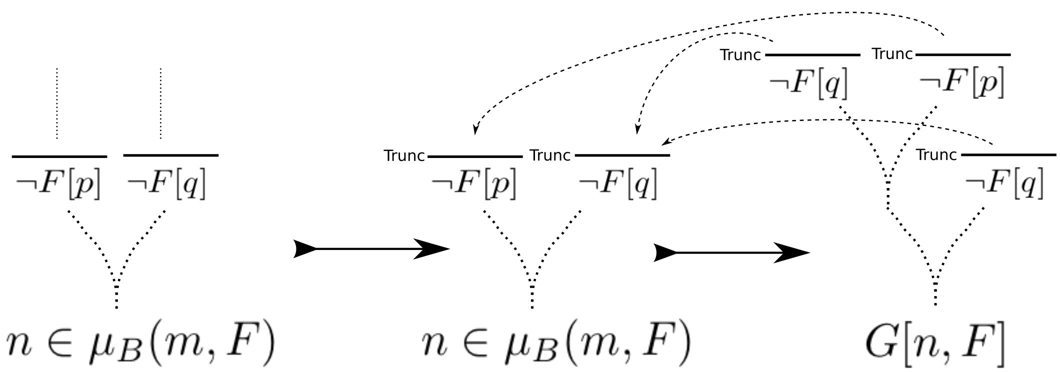

Fig. 1. Example Transformation

We cannot be finished, because we have thrown away everything above a widely branching inference in the original derivation. In order to recover it, we must provide, for each truncated inference appearing in our derivation of $G[n, F]$, not only a truncated inference from the source derivation, but also a new Ω rule which will provide, for each possible premise d_i, a new derivation $\mathcal{F}(d_i)$ in such a way that $\{\mathcal{F}(d_i)\}_\iota$ are valid premises for the widely branching inference.

In order to keep all this information in one place, our Ω rule for $n \in \mu_B(m, F)$ will branch, not over derivations, but over *sequences* of derivations. Given a derivation d of $n \in \mu_B(m, F)$, we divide this derivation into pieces by chopping it at each introduction of some $\neg F[k]$. We then build up a new derivation coinductively; the bottommost piece, d_0, is replaced by some $\mathcal{F}(\langle d_0\rangle)$. Each truncated inference θ appearing in $\mathcal{F}(d_0)$ is traced to some truncated inference in d_0, which in turn is traced to some introduction of $\neg F[k]$ in d using an inference rule $\mathcal{I}$. This introduction rule, whatever it is, has some list of premises $\{d_\iota\}$; for each d_ι there is an inference $\mathcal{F}(\langle d_0, d_\iota\rangle)$ which extends $\mathcal{F}(\langle d_0\rangle)$ at θ. By replacing θ in $\mathcal{F}(\langle d_0\rangle)$ with $\mathcal{I}$, taking for each premises ι the extension in $\mathcal{F}(\langle d_0, d_\iota\rangle)$, we obtain a new valid derivation. We then have new truncated inferences which first appeared in $\mathcal{F}(\langle d_0, d_\iota\rangle)$, and the process repeats.

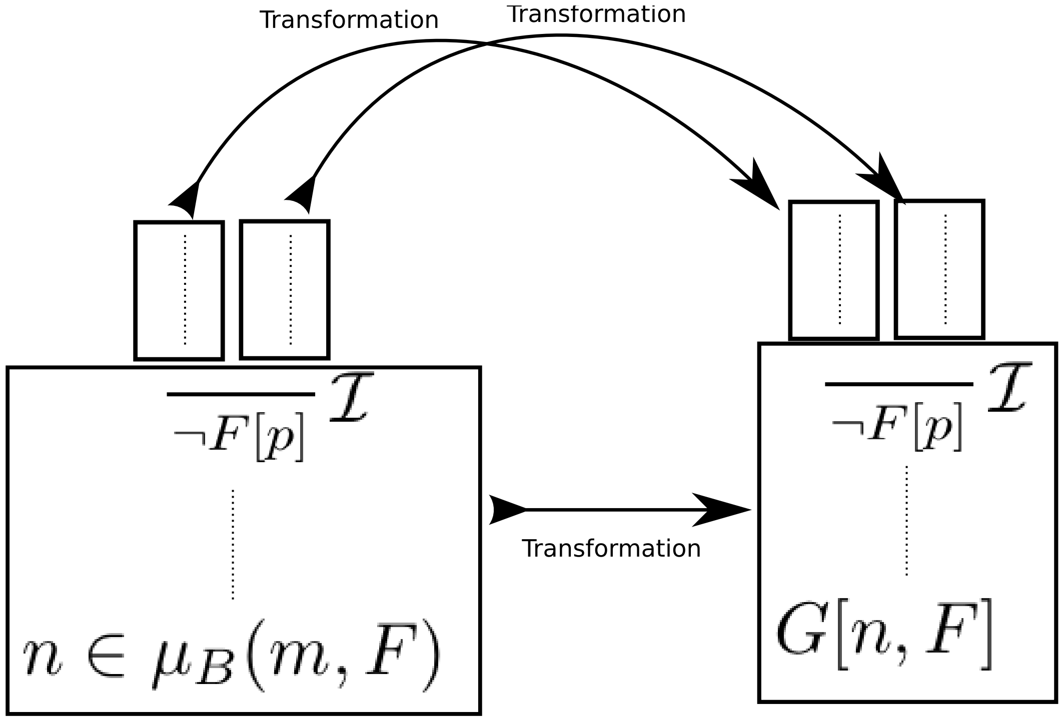

Fig. 2. A derivation is divided into segments, and (the corresponding portion of) the transformation is applied to each segment in turn.

We may formulate this procedure as a game with two players, a Prover and a Transformer. Prover plays first, and must play a derivation of from our system of small proofs augmented by truncated inferences (which we wall call a truncated proof system). Transformer must play a derivation with appropriate endsequent from the same system (actually, transformer is given a bit more flexibility, for instance, being allowed to use the cut rule), with the additional property that, for each truncated inference in this derivation, transformer must name a source callback inference in Prover's play. Prover then chooses some truncated inference in Transformer's play, and plays this truncated inference together with a new derivation. From here, play continues alternating these last two steps. Transformer wins as long as it is possible to provide derivations with the appropriate endsequent relative to what Prover offers (and an additional condition to be described shortly). The Ω rule is simply an encoding of a winning strategy for Transformer. (The ordinary Ω rule may be viewed as the two step version of this game, where Prover is not permitted an additional play after Transformer has gone once.) Any derivation gives a collection of strategies for Prover, and applying the transformation to some derivation is the result of knitting together the results given by Transformer against all the strategies for Prover offered by the derivation.

Two points must be made about this procedure. First, it is convenient in the description of cut-elimination to take the view that Transformer's plays (that is, the premises of the Ω rule) are not merely the portion of the derivation to be placed above truncated rules, but the entire derivation below that point as well. That is, $\mathcal{F}(\sigma^\frown\langle d_n\rangle, \tau^\frown\langle\theta\rangle)$ should be an extension of $\mathcal{F}(\sigma,\tau)$ in which the only change is that the truncated inference θ, which had no premise in $\mathcal{F}(\sigma,\tau)$, is required to have a single premise with appropriate endsequent (based on d_n) in $\mathcal{F}(\sigma^\frown\langle d_n\rangle, \tau^\frown\langle\theta\rangle)$. These truncated rules with an additional premise will be called callback inferences, since they represent the point at which Transformer's play has to make reference to the content omitted in Prover's play.

The second point is that truncated inferences appearing in $\mathcal{F}(\sigma,\tau)$ may have their source in any inference in σ, not just the most recent one. This is necessary, since the cut-elimination process will cause this situation to occur. However this introduces a concern about well-foundedness; we wish to have the property that whenever d is a well-founded derivation, the result of applying the transformation to it is also well-founded. In order to preserve this, we must specify additional conditions on infinite play; if Prover's plays are given by the infinite sequences σ, τ and the τ_i are all selected from the newly extended part of Transformer's play, Transformer loses if there is some σ_i such that infinitely many τ_j belong to σ_i. In any other infinite play, Prover loses. (A well-foundedness criterion of some sort is to be expected, since we are producing an analysis of a system stronger than Π^1_1-comprehension. It is not hard to show that a transformation with this property maps well-founded derivations to well-founded ones.) Our Ω rule must remain a winning strategy for this clarified version of the game.

Given this Ω rule, the remainder of our proof is not so difficult. Such Ω rules are considered an additional type of "small" inference, and may appear in derivations of $n \in \mu_A$, which then has an ordinary Ω rule.

3 Transformations

3.1 Proof System

We first need a general notion of a proof system, which we take almost verbatim from [13]. In the following, we assume we have already fixed some suitable language, and are working with the formulas of this language.

Definition 3.1. A *sequent* is a finite set of formulas.

A *proof system* consists of a set of formal *inference symbols* (generally denoted by the variable $\mathcal{I}$), and, for each inference symbol:

- A (possibly infinite) set $|\mathcal{I}|$ called its arity
- A sequent $\Delta(\mathcal{I})$
- For each $\iota \in |\mathcal{I}|$, a sequent $\Delta_\iota(\mathcal{I})$
- A set $Eig(\mathcal{I})$ which is either empty of a singleton $\{x\}$ where x is a variable not in $FV(\Delta(\mathcal{I}))$ (in this case we call x the *eigenvariable* of $\mathcal{I}$)

When we say that a proof system contains an inference rule
$\mathcal{I}\,\dfrac{\cdots\Delta_\iota\cdots(\iota\in I)}{\Delta}\,!u!$
we are declaring $\mathcal{I}$ to be an inference symbol with arity I, $\Delta(\mathcal{I}) = \Delta$, $\Delta_\iota(\mathcal{I}) = \Delta_\iota$, and $Eig(\mathcal{I}) = \{u\}$ (or $\emptyset$ if u is omitted). When the arity is finite, we typically list all the premises explicitly.

Definition 3.2. The *derivations* d of a proof system and the *end sequent* $\Gamma(d)$ are defined inductively. If, for each $\iota \in |\mathcal{I}|$, d_ι is a derivation and setting $\Gamma := \Delta(\mathcal{I}) \cup \bigcup_{\iota\in|\mathcal{I}|} \Gamma(d_\iota) \setminus \Delta_\iota(\mathcal{I})$, $Eig(\mathcal{I}) \cap FV(\Gamma) = \emptyset$ then $d := \mathcal{I}(d_\iota)_{\iota\in|\mathcal{I}|}$ is a derivation with $\Gamma(d) := \Gamma$.

If d is a derivation and $\Gamma(d) \subseteq \Gamma$ then we write $d \vdash \Gamma$.

Definition 3.3. An expression of the form $\lambda x.F$ is called a predicate, and denoted $\mathcal{F}$. We write $\mathcal{F}[t] := F(x/t)$.

3.2 Augmented and Truncated Derivations

We define proof systems with additional rules which serve to mark places where a derivation has been cut off. The rule $Trunc_{\Gamma\mapsto\Gamma,\Delta}$ indicates a point where the derivation has been truncated below an inference rule $\mathcal{I}$ with $\Delta(\mathcal{I}) = \Gamma$ and $\bigcup_{\iota\in|\mathcal{I}|} \Gamma(d_\iota) \setminus \Delta_\iota(\mathcal{I}) = \Delta$.

A $CB_{\Upsilon\mapsto\Delta}$ inference indicates a point where every branch besides the branch ι of some inference rule $\mathcal{I}$ has been cut off, $\Gamma(d_\iota) = \Upsilon$ and $\Delta(\mathcal{I}) \cup \Gamma(d_\iota) \setminus \Delta_\iota(\mathcal{I}) = \Delta$.

Definition 3.4. Let $\mathcal{P}$ be a proof system. We define *truncated* $\mathcal{P}$ to consist of $\mathcal{P}$ together with inference rules

$$Trunc_{\Gamma\mapsto\Gamma,\Delta}\,\frac{}{\Gamma,\Delta}$$

We define *augmented* $\mathcal{P}$ to consist of truncated $\mathcal{P}$ together with inference rules

$$CB_{\Upsilon\mapsto\Delta}\ \frac{\Upsilon}{\Delta}$$

We define $\Theta(d)$ to be the set of instances of $Trunc$ inferences appearing in d.

If θ is a truncated inference $Trunc_{\Gamma\mapsto\Gamma,\Delta}$, we set $In(\theta) := \Gamma$ and $Out(\theta) := \Delta$.

Note that Θ picks out instances, so it distinguishes two occurrences of the inference rule in different places, even if they have identical parameters.

We will want to be able to talk about systems such as truncated $\mathcal{P}$ where $\mathcal{P}$ is itself augmented $\mathcal{Q}$; when we speak of truncated inferences in a derivation in augmented $\mathcal{P}$, or refer to $\Theta(d)$, we mean to include only those inferences not belonging to $\mathcal{P}$. That is, augmenting and truncating give *disjoint* unions.

Definition 3.5. We define the *exploded derivations* of $\mathcal{P}$ over $\mathcal{Q}$ by induction:

- If d is a derivation in truncated $\mathcal{Q}$ and $\mathcal{I}, E$ are functions on $\Theta(d)$ such that $\mathcal{I}(\theta)$ is an inference rule from $\mathcal{P}$, $In(\theta) = \bigcup_\iota \Gamma(E(\theta,\iota)) \setminus \Delta_\iota(\mathcal{I}(\theta))$, $Out(\theta) = \Delta(\mathcal{I}(\theta))$, and each $E(\theta,\iota)$ is an exploded derivation then $\langle d, \mathcal{I}, E\rangle$ is an exploded derivation with endsequent $\Gamma(d)$

We denote the endsequent of an exploded derivation E by $\Gamma(E)$. If $E = \langle d, \mathcal{I}, E\rangle$ is an exploded derivation, we set $E_0 := d$ and call this the *main part* of the exploded derivation.

Definition 3.6. If $\langle d, \mathcal{I}, E\rangle$ is an exploded derivation, the *unexplosion* $\mathbb{U}(\langle d, \mathcal{I}, E\rangle)$ is given by main induction on E and a side induction on d:

- If d is a $Trunc$ inference,

$$\mathbb{U}(\langle Trunc, \mathcal{I}, E\rangle) := \mathcal{I}(\theta)\{\mathbb{U}(E(\theta,\iota))\}_{\iota\in|\mathcal{I}(\theta)|}$$

- Otherwise,

$$\mathbb{U}(\langle \mathcal{J}\{d_\iota\}, \mathcal{I}, E\rangle) := \mathcal{J}\{\mathbb{U}(\langle d_\iota, \mathcal{I}\restriction\Theta(d_\iota), E\restriction\Theta(d_\iota)\rangle)\}_{\iota\in|\mathcal{J}|}$$

Definition 3.7. If $\mathcal{P}, \mathcal{Q}$ are proof systems, Δ a sequent, and $\mathcal{F}$ a predicate, we define the *explosion* $\mathbb{E}_{\mathcal{Q}}(d)$ of a derivation d in $\mathcal{P}$ by:

- If $d = \mathcal{I}\{d_\iota\}$ where $\mathcal{I}$ is not an inference of $\mathcal{Q}$,

$$\mathbb{E}_{\mathcal{Q}}(d) := \langle Trunc_{\Gamma(d)\setminus\Delta(\mathcal{I})\to\Gamma(d)}, \theta\mapsto\mathcal{I}, (\theta,\iota)\mapsto\mathbb{E}_{\mathcal{Q}}(d_\iota)\rangle$$

- Otherwise $d = \mathcal{I}\{d_\iota\}$ where $\mathcal{I}$ is an inference of $\mathcal{Q}$ and set, for each $\iota\in|\mathcal{I}|$, $\langle d'_\iota, \mathcal{I}_\iota, E_\iota\rangle := \mathbb{E}_{\mathcal{Q}}(d)$, and then

$$\mathbb{E}_{\mathcal{Q}}(d) := \langle \mathcal{I}\{d'_\iota\}, \bigcup \mathcal{I}_\iota, \bigcup E_\iota\rangle$$

Lemma 3.1. For any $\mathcal{Q}$, $\mathbb{U}(\mathbb{E}_{\mathcal{Q}}(d)) = d$.

Definition 3.8. Let d, d' be derivations in augmented $\mathcal{P}$ such that d and d' are identical except that there exist some $Trunc_{\Gamma \mapsto \Gamma, \Delta}$ inference in d, but at the corresponding place in d', there is a $\theta = CB_{\Upsilon \mapsto \Gamma, \Delta}$ inference. We say d' *narrowly extends* d, and write $d' \setminus d$ for the derivation which is the premise of the callback inference θ in d'. We call Υ the *key sequent* of this extension.

Definition 3.9. Let $\mathcal{P}, \mathcal{Q}$ be proof systems, and let sequences of equal length σ_0, τ_0 be given. We say $\{d_{\sigma,\tau}\}_{\sigma \supseteq \sigma_0, \tau \supseteq_0 \tau}$ together with supplementary functions $\Lambda_{\sigma,\tau}$ is a *transformation* from Γ out of $\mathcal{P}$ over some restricted set of formulas $\mathcal{F}$ (in a proof system $\mathcal{Q}$) with endsequent Σ and root σ_0, τ_0 if the following holds:

- For every derivation d of Γ, Υ in truncated $\mathcal{P}$ with $\Upsilon \subseteq \mathcal{F}$, $d_{\sigma_0^\frown \langle d \rangle, \tau_0}$ is a proof of Σ, Υ in truncated $\mathcal{Q}$, $\Lambda_{\sigma_0^\frown \langle d \rangle, \tau_0} : \Theta(d_{\sigma_0^\frown \langle d \rangle, \tau_0}) \to \Theta(d) \cup \bigcup_{i < length(\sigma)} \Theta(\sigma_i)$, and for each θ, $Out(\theta) = Out(\Lambda_{\sigma_0^\frown \langle d \rangle, \tau_0}(\theta))$
- If $d_{\sigma,\tau}$ is defined, $\theta \in \Theta(d_{\sigma,\tau})$, and $d \vDash In(\Lambda_{\sigma,\tau}(\theta)), \Upsilon$ in truncated $\mathcal{P}$ with $\Upsilon \subseteq \mathcal{F}$ then $d_{\sigma^\frown \langle d \rangle, \tau^\frown \langle \theta \rangle}$ is a proof in augmented $\mathcal{Q}$ narrowly extending $d_{\sigma,\tau}$ at θ and $d_{\sigma^\frown \langle d \rangle, \tau^\frown \langle \theta \rangle}$ is a proof in truncated $\mathcal{Q}$ with key sequent $In(\theta), \Upsilon$. Furthermore, $\Lambda_{\sigma^\frown \langle d \rangle, \tau^\frown \langle \theta \rangle}$ has range in $\bigcup \Theta(\sigma_i)$ and agrees with $\Lambda_{\sigma^\frown \langle d \rangle, \tau}$ on elements in their shared domain.

If $T = \{d_{\sigma,\tau}\}_{\sigma \supseteq \sigma_0, \tau \supseteq \tau_0}$ is a transformation and $\sigma' \supseteq \sigma_0, \tau' \supseteq \tau_0$ are such that $d_{\sigma',\tau'}$ is defined, we write $T \upharpoonright \sigma', \tau'$ for the transformation $\{d_{\sigma,\tau}\}_{\sigma \supseteq \sigma', \tau \supseteq \tau'}$.

Let d be given and let σ_0, τ_0 be given with d an element of σ. We define the *d-well-founded transformations* inductively:

- If there is no $\sigma \supseteq \sigma_0, \tau \supseteq \tau_0$, $\theta \in \Theta(d_{\sigma,\tau})$ such that $\Lambda_{\sigma,\tau}(\theta) \in \Theta(d)$ then $\{d_{\sigma,\tau}\}_{\sigma \supseteq \sigma_0, \tau \supseteq \tau_0}$ is d-well-founded. We call such transformations *d-void*.
- If for every d', θ, $T \upharpoonright \sigma_0^\frown \langle d' \rangle, \tau_0^\frown \langle \theta \rangle$ is d-well-founded then so is T

We say $T = \{d_{\sigma,\tau}\}$ is *well-founded* if for every σ, τ and every d such that $d_{\sigma^\frown \langle d \rangle, \tau}$ is defined, $T \upharpoonright \sigma^\frown \langle d \rangle, \tau$ is d-well-founded.

A transformation should, as the name suggests, give a way of transforming a derivation of Γ, Υ into a derivation of Σ, Υ. In order to get the right inductive hypothesis, we need to first show how to apply a transformation to an exploded derivation.

Lemma 3.2. Let $E = \langle d_0, \mathcal{I}, E_0 \rangle$ be an exploded derivation over $\mathcal{P}$ with endsequent Γ, Υ and let $T = \{d_{\sigma,\tau}\}_{\sigma \supseteq \sigma_0, \tau \supseteq \tau_0}$ be a well-founded transformation from Γ out of $\mathcal{P}$ over some $\mathcal{F} \supseteq \Upsilon$ with endsequent Σ. Then there is a derivation d^* with endsequent Σ, Υ and a function $\Lambda : \Theta(d^*) \to \bigcup \Theta((\sigma_0)_i)$.

Proof. The proof is by main induction on E and side induction on T. Let $E = \langle d_0, \mathcal{I}, E_0 \rangle$ be given. Then by induction, we produce from any d_0-wellfounded transformation T a d_0-void transformation T'. If T is d_0-void then $T' = T$. Otherwise, by the side IH, for each d', θ, $T \upharpoonright \sigma_0^\frown \langle d' \rangle, \tau_0^\frown \langle \theta \rangle$ is d_0-wellfounded. For each d let $d'_{\sigma_0^\frown \langle d \rangle, \tau_0}$ be the result of replacing each $\theta \in \Theta(d_{\sigma_0^\frown \langle d \rangle, \tau_0})$ such that $\Lambda_{\sigma_0^\frown \langle d \rangle, \tau_0}(\theta) \in \Theta(d_0)$, with $\mathcal{I}(\Lambda_{\sigma_0^\frown \langle d \rangle, \tau_0}(\theta))$ with the premise ι given by T' applied to $E_0(\Lambda_{\sigma_0^\frown \langle d \rangle, \tau_0}(\theta), \iota)$; this application exists by the main IH.

Then $d^* := d'_{\sigma_0^\frown \langle d_0 \rangle, \tau_0}$ and $\Lambda := \Lambda_{\sigma_0^\frown \langle d_0 \rangle, \tau_0} \upharpoonright \Theta(d^*)$ witness the theorem. □

Theorem 3.1. If T is a transformation out of $\mathcal{Q}$ from Γ over $\mathcal{F}$ with endsequent Σ and $d \vDash \Gamma, \Upsilon$ for some $\Upsilon \subseteq \mathcal{F}$ then there is a derivation $T(d)$ of Σ, Υ.

Proof. Apply the preceding lemma to $\mathbb{E}_{\mathcal{Q}}(d)$. □

Lemma 3.3. Let T be a wellfounded transformation, let $\{\mathcal{O}_i\}$ be a set of operators on derivations, all with the same domain, and for each $\mathcal{O}_i$, let $\Lambda_{\mathcal{O}_i}$ be a function with the properties that:

- Each $\mathcal{O}_i$ takes wellfounded derivations to wellfounded derivations
- Each $\mathcal{O}_i$ preserves narrow extensions; that is, if d' narrowly extends d then $\mathcal{O}_i(d')$ narrowly extends $\mathcal{O}_i(d)$
- For every d in the domain of $\mathcal{O}_i$, $\Lambda_{\mathcal{O}_i}(d) : \Theta(\mathcal{O}_i(d)) \to \Theta(d)$ with the property that $Out(\theta) = Out(\Lambda_{\mathcal{O}_i}(d)(\theta))$ and if $d' \vDash In(\Lambda_{\mathcal{O}}(d)(\theta)), \Upsilon$ belongs to the domain then there is an operator $\mathcal{O}_j$ such that $\mathcal{O}_j(d') \vDash In(\theta), \Upsilon$.

Then each $\mathcal{O}_i$ extends to an operator on wellfounded transformations, $T \mapsto \mathcal{O}_i \circ T$, with appropriate domain and range with the property that

$$(\mathcal{O}_i \circ T)(d) = \mathcal{O}_i(T(d))$$

for any derivation d.

Proof. Follows immediately by applying operators pointwise, using $\Lambda_{\mathcal{O}_i}$ to define $\mathcal{O}_i(\Lambda)$. □

We call such a system of such operators *uniform*.

4 The System μ_2

4.1 Language

Definition 4.1. If $A(X,x)$ is a formula, we write $A(X)$ for $\{x \mid A(X,x)\}$; in particular, $A(X) \subseteq X$ means $\forall x(A(X,x) \rightarrow x \in X)$.

As we define our system, we also assign depths to formulas. Depths will be ordinals $\leq \omega+\omega$, although we will immediately restrict ourselves to $\omega+2$. (The use of the ordinal $\omega+\omega$ is somewhat artificial; we have ω levels corresponding to finitely many iterated inductive definitions, and then three levels above, corresponding to the inaccessible, the negated inaccessible, and an admissible above the inaccessible. The names $< I$, I, and $I+1$ might convey this more clearly.)

Definition 4.2. The language of $\mathcal{L}_{\mu_2}$ is defined as follows:

- 0 is a constant symbol
- S is a unary function constant symbol
- There are infinitely many symbols for variables
- For each n-ary primitive recursive relation, including $=$ and $\leq$, there is an n-ary predicate constant symbol R
- The logical symbols are $\neg, \wedge, \vee, \forall, \exists$
- If $A(x,X)$ contains no other free variables and contains X positively then $\mu xX.A(x,X)$ is a unary predicate symbol
- If $B(y,Y,Z)$ contains Y positively and Z negatively and $A(x,X,Z)$ contains X positively and Z negatively, and A and B have finite depth then $\mu xX.A(x,X,\mu yY.B(y,Y,X))$ is a unary predicate symbol; we call this a predicate of *inaccessible type*

The terms are given by:

- 0 is a term
- If t is a term then St is a term
- Each variable is a term

The formulas are given by:

- If R is a symbol for an n-ary primitive recursive relation and for each $i \leq n$, t_i is a term, then $Rt_1 \ldots t_n$ is an atomic formula of depth n for any $n \geq 0$
- If $A(x,X)$ has depth n and t is a term then $t \in \mu xX.A(x,X)$ is an atomic formula of depth n
- If t is a term then $t \in \mu xX.A(x,X,\mu yY.B(y,Y,X))$ is an atomic formula of depth ω

- If A is an atomic formula of depth n, $\neg A$ is a formula of depth $n+1$
- If A_0 and A_1 are formulas of depth n then $A_0 \wedge A_1$ and $A_0 \vee A_1$ are formulas of depth n
- If x is a variable and A a formula of depth n then $\forall x A$ and $\exists x A$ are formulas of depth n

If $n < \omega$ then $\mathcal{L}_{ID_n}$ is the restriction to formulas of depth n. The depth of a formula, $dp(A)$, is the least $n \geq 0$ such that A has depth n.

If $dp(A) \geq \omega + 1$ then we call $\mu x X.A(x, X)$, and any formula containing it, *large*.

Our theory will effectively restrict consideration to formulas of depth at most $\omega+2$. Note that all formulas of higher depth are large. The restriction is somewhat artificial, since we have to "throttle" the formation rule for μ-expressions, but the alternative would be analyzing a stronger system corresponding to an inaccessible with infinitely many admissibles above it. (This phenomenon has been observed before, for instance in [16], where the addition of a constructor corresponding to an inaccessible immediately pushes the system up to infinitely many admissibles beyond it due to the presence of other constructors.)

Definition 4.3. $FV(\phi)$ denotes the set of free variables of ϕ, and ϕ is closed if $FV(\phi) = \emptyset$. Here ϕ may be a formula, a term, or a sequent.

If A is not atomic, $\neg A$ indicates the negation of A in negation normal form as given by de Morgan's laws.

The rank $rk(A)$ of a formula is defined by:

- $rk(A) := 0$ if A is atomic
- $rk(\neg A) := rk(A)$
- $rk(A \wedge B) = rk(A \vee B) := \max\{rk(A), rk(B)\} + 1$
- $rk(\forall x A) = rk(\exists x A) := rk(A) + 1$

$A(x/t)$ means the result of substituting t for every free occurrence of x in A (renaming bound variables if necessary). When x is clear, we just write $A(t)$.

Definition 4.4. We define the true primitive recursive formulas to be those closed primitive recursive atomic formulas and negations of atomic formulas which are true in the standard interpretation.

The system μ_2 contains the following inference symbols:

$Ax_\Delta \ \dfrac{}{\Delta}$
where Δ contains a true primitive recursive formula or a pair $t \in \mu x X.A(x, X), n \notin$

$\mu x X.A(x,X)$

$$\bigwedge_{A_0\wedge A_1}\frac{A_0 \qquad A_1}{A_0\wedge A_1} \qquad \bigwedge^{i}_{A_0\vee A_1}\frac{A_i}{A_0\vee A_1}$$

$i\in\{0,1\}$

$$\bigwedge^{y}_{\forall x A}\frac{A(y)}{\forall x A}\,!x! \qquad \bigvee^{t}_{\exists x A}\frac{A(t)}{\exists x A}$$

$$Cut_C\,\frac{C \qquad \neg C}{\emptyset} \qquad Ind^{t}_{\mathcal{F}}\,\frac{}{\neg\mathcal{F}[0],\neg\forall x(\mathcal{F}[x]\to\mathcal{F}[Sx]),\mathcal{F}[t]}$$

where C is not large

$$Cl_{t\in\mu x X.A(x,X)}\,\frac{A(t,\mu x X.A(x,X))}{t\in\mu x X.A(x,X)}$$

$$Ind^{\mu x X.A(x,X),t}_{\mathcal{F}}\,\frac{}{\neg(A(\mathcal{F})\subseteq\mathcal{F}),t\notin\mu x X.A(x,X),\mathcal{F}[t]}$$

We say a derivation d belongs to ID_n if every formula in every endsequent in d belongs to $\mathcal{L}_{ID_n}$.

4.2 Infinitary Derivations

We define an infinitary system μ_2^∞; its language is the same language $\mathcal{L}_{\mu_2}$, but only closed formulas are permitted. This definition will require that a number of weaker systems be defined along the way.

The following, which we will call ID_0^∞, will be the basis for all the systems we need. Roughly, it is the standard infinitary system for Peano Arithmetic plus a closure rule—but not an induction rule—for $\mu x X.A(x,X)$ of depth 0.

Definition 4.5. $Ax_\Delta\,\frac{}{\Delta}$

where Δ contains a true primitive recursive formula

$$\bigwedge_{A_0 \wedge A_1} \frac{A_0 \qquad A_1}{A_0 \wedge A_1} \qquad \bigvee^{i}_{A_0 \vee A_1} \frac{A_i}{A_0 \vee A_1}$$

$i \in \{0, 1\}$

$$\bigwedge_{\forall x A} \frac{\cdots A(i) \cdots \qquad (i \in \mathbb{N})}{\forall x A} \qquad \bigvee^{n}_{\exists x A} \frac{A(n)}{\exists x A}$$

$n \in \mathbb{N}$

$$Cl_{n \in \mu x X.A(x,X)} \frac{A(n, \mu x X.A(x, X))}{n \in \mu x X.A(x, X)} \qquad Cut_C \frac{C \qquad \neg C}{\emptyset}$$

and all formulas have depth 0.

Definition 4.6. If q is a proof and Γ a sequent, $\Delta^{\Gamma}_{q} := \Gamma(q) \setminus \Gamma$.

The systems ID^{∞}_{n+1} are defined inductively; as the name suggests, they are essentially the infinitary systems from [13].

Definition 4.7. Given ID^{∞}_{n}, the language of the system ID^{∞}_{n+1} is $\mathcal{L}_{ID_n}$—that is, formulas with depth $\leq n+1$, and consists of the rules of ID^{∞}_{n} together with

$$Ax_{\Delta} \frac{}{\Delta}$$

where Δ contains $n \in \mu x X.A(x, X), n \notin \mu x X.A(x, X)$ with $dp(\mu x X.A(x, X)) < n+1$

$$\Omega_{k \notin \mu x X.A(x,X)} \frac{k \in \mu x X.A(x, X) \qquad \ldots \Delta^{k \in \mu x X.A(x,X)}_{q} \ldots \qquad (q \in |k \in \mu x X.A(x, X)|)}{\emptyset}$$

where $|k \in \mu x X.A(x, X)|$ is the set of cut-free proofs of $ID^{\infty}_{dp(k \in \mu x X.A(x,X))}$ and $dp(\mu x X.A(x, X)) \leq n$, and $\Delta_q(\Omega_{k \notin \mu x X.A(x,X)}) := \Upsilon$ where $q \vdash k \in \mu x X.A(x, X), \Upsilon$

Note that the premise of the Ω rule d defines a function taking proofs of $k \in \mu x X.A(x, X)$ to proofs of $\Gamma(d)$.

Definition 4.8. Next we define a system μ^{∞}_{ω}, which extends the union of ID^{∞}_{n} with the closure rule for predicates of inaccessible type.

Note that this doesn't add any derivations—there's no way to introduce $A(n, \mu_A)$ since there's no way to introduce $n \notin \mu_B(\mu_A)$. We're including the rule so that it will be present in the extensions we need.

Definition 4.9. The system μ^{∞}_{I} extends μ^{∞}_{ω} by the rule

$$\neg_{n \notin \mu x X.A(x,X)} \frac{n \in \mu x X.A(x, X, \mu_1, \ldots, \mu_k) \qquad \ldots d_{\sigma,\tau} \ldots}{\emptyset}$$

where $\mu_1, \ldots, \mu_k$ are predicates of inaccessible type appearing negatively in A, no other predicates of inaccessible type appear in A, and for every $\mathcal{F}_1, \ldots, \mathcal{F}_k$, the premises include a well-founded transformation from $n \in \mu x X.A(x, X, \mathcal{F}_1, \ldots, \mathcal{F}_k)$ out of the cut-free part of μ^∞_ω over $\mu x X.A(x, X, \mathcal{F}_1, \ldots, \mathcal{F}_k)$ positive formulas.

Now we can define our final system:

Definition 4.10. The system μ^∞ consists of μ^∞_I plus the rules

$$Ax_\Delta \frac{}{\Delta}$$

where Δ contains $n \in \mu x X.A(x, X), n \notin \mu x X.A(x, X)$ and $\mu x X.A(x, X)$ has inaccessible type

$$\Omega_{n \notin \mu x X.A(x,X)} \frac{n \in \mu x X.A(x, X) \qquad \ldots d_{\sigma,\tau} \ldots}{\emptyset}$$

where the premises range over cut-free proofs of μ^∞_I

Note that none of these systems allow cut rules over large formulas.

Definition 4.11. Given a system $\mathcal{P}$, the *augmentations* of $\mathcal{P}$ are given inductively: $\mathcal{P}$ is an augmentation of $\mathcal{P}$, and if $\mathcal{Q}$ is an augmentation of $\mathcal{P}$ then so is augmented $\mathcal{Q}$.

Definition 4.12. We define $c{-}rk(d)$, the cut-rank of d, inductively as follows:

$$c - rk(d) = \max\{c - rk(d_\iota) \mid \iota \in |\mathcal{I}|\}$$

unless $\mathcal{I} = Cut_C$

$$c - rk(Cut_C(d_0, d_1)) = \max\{c - rk(d_0), c - rk(d_1), rk(C) + 1\}$$

5 Embedding

Definition 5.1. A derivation in μ is closed if every number variable occurring free is the eigenvariable of an inference below that occurrence. In particular, $FV(\Gamma(h)) = \emptyset$ if h is closed.

We will define a function taking closed proofs in μ_2 to proofs in μ^∞. The hard part will be the induction axioms, which will be embedded as Ω rules. Most of the work is defining the functions used to make these Ω rules.

Definition 5.2. Let $\mathbf{d}_{\mathcal{F},\neg\mathcal{F}}$ be the canonical derivation of $\mathcal{F}, \neg\mathcal{F}$.

If $d \vdash A(n, \mathcal{F})$ then $\mathbf{e}^{n}_{\mathcal{F},A}(d)$ is the derivation

$$\dfrac{\dfrac{\begin{array}{c} d \\ \vdots \\ A(n,\mathcal{F}) \end{array} \qquad \begin{array}{c} \mathbf{d}_{\mathcal{F}[n],\neg\mathcal{F}[n]} \\ \vdots \\ \mathcal{F}[n],\neg\mathcal{F}[n] \end{array}}{\mathcal{F}[n], A(n,\mathcal{F}) \wedge \neg\mathcal{F}[n]}}{\mathcal{F}[n], \neg(A(\mathcal{F}) \subseteq \mathcal{F})}$$

or symbolically

$$\bigvee^{n}_{\neg(A(\mathcal{F})\subseteq\mathcal{F})} \ \bigwedge_{A(n,\mathcal{F})\wedge\neg\mathcal{F}[n]} d\mathbf{d}_{\neg(\mathcal{F}[n]),\mathcal{F}[n]}$$

Lemma 5.1. There is a function $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}$ such that if $dp(\mu xX.A(x,X)) < \omega$ and $d \vdash \Pi(\mu xX.A(x,X)), \Sigma$ is a cut-free proof in $ID^{\infty}_{dp(\mathcal{G})}$ then

$$\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d) \vdash \Pi(\mathcal{F}), \neg(A(\mathcal{F}) \subseteq \mathcal{F}), \Sigma$$

is a proof in μ^{∞}.

Proof. By induction on d. We simply proceed up through the proof, adding to Π as we encounter subformulas or new formulas produced by closures rules. A typical case is

$$\dfrac{B_0(\mu xX.A(x,X)) \qquad B_1(\mu xX.A(x,x))}{B_0(\mu xX.A(x,X)) \wedge B_1(\mu xX.A(x,x))} \quad \mapsto \quad \dfrac{B_0(\mathcal{F}) \qquad B_1(\mathcal{F})}{B_0(\mathcal{F}) \wedge B_1(\mathcal{F})}$$

where $B_0 \wedge B_1$ belongs to Π.

The only difficult case is the closure rule, which we handle with the help of $\mathbf{e}$:

$$\dfrac{A(n, \mu xX.A(x,X))}{n \in \mu xX.A(x,X)} \quad \mapsto \quad \dfrac{\dfrac{A(n,\mathcal{F}) \qquad \begin{array}{c} \mathbf{d}_{\mathcal{F}[n],\neg\mathcal{F}[n]} \\ \vdots \\ \mathcal{F}[n],\neg\mathcal{F}[n] \end{array}}{\mathcal{F}[n], A(n,\mathcal{F}) \wedge \neg\mathcal{F}[n]}}{\mathcal{F}[n], \neg(A(\mathcal{F}) \subseteq \mathcal{F})}$$

Importantly, we never encounter $n \notin \mu xX.A(x,X)$ anywhere; in particular, we do not have to deal with the axiom $Ax_{n\in\mu xX.A(x,X), n\notin\mu xX.A(x,X)}$.

The full definition is given by

$$\mathcal{SUB}^{\Pi}_{n\in\mu xX.A(x,X),\mathcal{F}}(\mathcal{I}(d_{\iota})_{\iota\in|\mathcal{I}|}) :=$$

$$\begin{cases} \mathbf{e}^n_{\mathcal{F},\mathcal{G}}(\mathcal{SUB}^{\Pi\cup\{\Delta_0(\mathcal{I})\}}_{\mu xX.A(x,X),\mathcal{F}}(d_0)) & \text{if } \mathcal{I} = Cl_{n\in\mu xX.A(x,X)} \\ & \text{ and } n \in \mu xX.A(x,X) \in \Pi \\ \mathcal{I}_{A(\mathcal{F})}\left(\mathcal{SUB}^{\Pi\cup\{\Delta_\iota(\mathcal{I})\}}_{\mu xX.A(x,X),\mathcal{F}}(d_\iota)_{\iota\in|\mathcal{I}|}\right) & \text{if } \mathcal{I} = \mathcal{I}_{B(\mu xX.A(x,X))} \\ & \text{ and } B(\mu xX.A(x,X)) \in \Pi \\ \mathcal{I}_A(\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d_\iota)_{\iota\in|\mathcal{I}|}) & \text{otherwise} \end{cases}$$

□

Lemma 5.2. Let $A(x,X)$ be a formula. Then there is an operator $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}$ such that if $d \vdash \Pi(\mu xX.A(x,X)), \Sigma$ is a cut-free proof in an augmentation of $\mu^\infty_{<I}$ then

$$\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d) \vdash \Pi(\mathcal{F}), \neg(A(\mathcal{F}) \subseteq \mathcal{F}), \Sigma$$

is a proof in the corresponding augmentation of μ^∞. Furthermore, this operator is uniform.

Proof. By induction on d. The proof is essentially the same as in the previous lemma, except that we add an additional case to handle $Trunc$ and CB inferences.

$$\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(\mathcal{I}(d_\iota)_{\iota\in|\mathcal{I}|}) :=$$

$$\begin{cases} \mathbf{e}^n_{\mathcal{F}}(\mathcal{SUB}^{\Pi\cup\{\Delta_0(\mathcal{I})\}}_{\mu xX.A(x,X),\mathcal{F}}(d_0)) & \text{if } \mathcal{I} = Cl_{n\in\mu xX.A(x,X)} \\ & \text{ and } n \in \mu xX.A(x,X) \in \Pi \\ Trunc_{F\mapsto\Pi(\mathcal{F}),\Sigma}\mathcal{SUB}^{\Pi}_{\mu xX.A(x,x),\mathcal{F}}(d_0) & \text{if } \mathcal{I} = Trunc_{F\mapsto\Pi(\mu xX.A(x,X)),\Sigma} \\ CB_{\Pi(\mathcal{F}),\Sigma}\mathcal{SUB}^{\Pi}_{\mu xX.A(x,x),\mathcal{F}}(d_0) & \text{if } \mathcal{I} = CB_{\Pi(\mu xX.A(x,X)),\Sigma} \\ \mathcal{I}_{A(\mathcal{F})}\left(\mathcal{SUB}^{\Pi\cup\{\Delta_\iota(\mathcal{I})\}}_{\mu xX.A(x,X),\mathcal{F}}(d_\iota)_{\iota\in|\mathcal{I}|}\right) & \text{if } \mathcal{I} = \mathcal{I}_{B(\mu xX.A(x,X))} \\ & \text{ and } B(\mu xX.A(x,X)) \in \Pi \\ \mathcal{I}_A(\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d_\iota)_{\iota\in|\mathcal{I}|}) & \text{otherwise} \end{cases}$$

□

Lemma 5.3. Let $A(x,X)$ be a formula. Then there is an operator $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}$, and a companion $\Lambda_{\sigma,\tau}$, giving a well-founded transformation.

Proof. By induction on the length of σ. $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X,\mathcal{G}_1,\mathcal{G}_k),\mathcal{F}}(\langle d_0\rangle, \langle\rangle)$ is just $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d_0)$ as given by the previous lemma. Given $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(\sigma,\tau)$, $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(\sigma^\frown\langle d\rangle, \tau^\frown\langle\theta\rangle)$ is given by replacing θ in $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(\sigma,\tau)$ with the derivation $\mathcal{SUB}^{\Pi'}_{\mu xX.A(x,X),\mathcal{F}}(d)$ where Π' is chosen to be the unique sequent such that θ was equal to $\mathcal{SUB}^{\Pi'}_{\mu xX.A(x,X),\mathcal{F}}(\Lambda_{\sigma,\tau}(\theta))$.

The function $\Lambda_{\sigma,\tau}$ is simply the association of each truncated inference in the range with the corresponding inference in the domain. $\square$

At last, we come to the key lemma:

Lemma 5.4. If $\mu x X.A(x,X)$ has inaccessible type, there is an operator $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}$ such that whenever d is a proof of $\Pi(\mu xX.A(x,X)),\Gamma$ in μ^{∞}_{I} then

$$\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d) \vdash \Pi(\mathcal{F}),\Gamma,\neg(A(\mathcal{F})\subseteq\mathcal{F})$$

Furthermore, $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}$ is uniform.

Proof. First, the simple cases are given by

$$\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(\mathcal{I}(d_\iota)_{\iota\in|\mathcal{I}|},F\}) :=$$

$$\begin{cases} \mathbf{e}^{n}_{\mathcal{F}}(\mathcal{SUB}^{\Pi\cup\{\Delta_0(\mathcal{I})\}}_{\mu xX.A(x,X),\mathcal{F}}(d_0)) & \text{if } \mathcal{I}=Cl_{n\in\mu xX.A(x,X)} \\ & \text{and } n\in\mu xX.A(x,X)\in\Pi \\ \mathcal{I}_{A(\mathcal{F})}\left(\mathcal{SUB}^{\Pi\cup\{\Delta_\iota(\mathcal{I})\}}_{\mu xX.A(x,X),\mathcal{F}}(d_\iota)_{\iota\in|\mathcal{I}|}\right) & \text{if } \mathcal{I}=\mathcal{I}_{A(\mu xX.A(x,X))} \\ & \text{and } A(\mu xX.A(x,X))\in\Pi \end{cases}$$

Next, consider $\neg_{n\notin\mu yY.B(y,Y,\mu xX.A(x,X),\mu_2,\ldots,\mu_k)}$ where $n\notin\mu yY.B(y,Y,\mu xX.A(x,X),\mu_2,\ldots,\mu_k)\in\Pi$. We use the abbreviations μ_A and $\mu_B(\mu_A)$ as in the introduction, and let T be the transformation formed by the premises. First, consider the simplest case, where $\mathcal{F}$ does not contain predicates of inaccessible type and $k=1$. Then we simply need to produce a function for an $\Omega_{n\notin\mu_B(\mathcal{F})}$ inference.

Since the premises give a transformation showing $n\in\mu_B(\mathcal{F})\mapsto\Pi(\mu xX.A(x,X)),\Gamma\setminus\{n\notin\mu yY.B(y,Y,\mu xX.A(x,X))\}$, also $\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}\circ d_{\sigma,\tau}$ gives a transformation T showing $n\in\mu_B(\mathcal{F})\mapsto\Pi(\mathcal{F}),\Gamma\setminus\{n\notin\mu yY.B(y,Y,\mathcal{F})\}$. Then we may assign to each $q\vdash n\in\mu_B(\mathcal{F}),\Upsilon$ the derivation

$$d_q := \mathbb{U}(T_*(\mathbb{E}_{\mu^{\infty}_{I}}(q)))$$

More generally, if predicates of inaccessible type occur in $\mathcal{F}$ or $k>1$ the same argument gives many transformations which collectively witness the corresponding $\neg$ inference.

In any other case, we do nothing:

$$\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(\mathcal{I}(d_\iota)) := \mathcal{I}_A(\mathcal{SUB}^{\Pi}_{\mu xX.A(x,X),\mathcal{F}}(d_\iota)_{\iota\in|\mathcal{I}|})$$

$\square$

Lemma 5.5. If h is a closed μ_2 derivation of Δ with $dp(\Delta) \leq \omega + 2$ then there is a μ^∞ derivation h^∞ so that $h^\infty \vdash_m \Gamma(h)$ for some finite m.

Proof. We define the $\cdot^\infty$ operation by induction on the proof h:

- $(\bigwedge^y_{\forall xA} d_0)^\infty := \bigwedge_{\forall xA}(d_0[n]^\infty)_{n\in\mathbb{N}}$
- $(Ind^0_{\mathcal{F}})^\infty := \mathbf{d}_{\mathcal{F}[0],\neg\mathcal{F}[0]}$
- $(Ind^{n+1}_{\mathcal{F}})^\infty := \bigvee^n_{\exists x(\mathcal{F}[x]\wedge\neg\mathcal{F}[Sx])} \bigwedge_{\mathcal{F}[n]\wedge\neg\mathcal{F}[Sn]}(Ind^n_{\mathcal{F}})^\infty \mathbf{d}_{\neg\mathcal{F}[Sn],\mathcal{F}[Sn]}$
- $(Ind^{\mu_A,n}_{\mathcal{F}})^\infty := \Omega_{n\notin\mu_A} Ax_{n\in\mu_A,n\notin\mu_A}\{\mathcal{SUB}^{n\in\mu_A}_{\mu_A,\mathcal{F}}(q)\}$ if $dp(\mu_A) < \omega$ or has inaccessible type
- $(Ind^{\mu_A,n}_{\mathcal{F}})^\infty := \neg_{n\notin\mu_A} Ax_{n\in\mu_A,n\notin\mu_A}\{\mathcal{SUB}^{n\in\mu_A(\mathcal{F}_1,\ldots,\mathcal{F}_k)}_{\mu_A(\mathcal{F}_1,\ldots,\mathcal{F}_k)),\mathcal{F}}(\sigma,\tau)\}_{\sigma,\tau,\mathcal{F}_1,\ldots,\mathcal{F}_k}$ if $dp(\mu_A) \geq \omega$, $\mu_A(\mu_1,\ldots,\mu_k)$ does not have inaccessible type, and the μ_i are all predicates of inaccessible type appearing in A
- Otherwise $(\mathcal{I}h_0\ldots h_{n-1})^\infty := \mathcal{I}h_0^\infty\ldots h_{n-1}^\infty$

□

6 Cut-Elimination

Definition 6.1. We say that A has $\bigwedge$-Form if it is either $A_0 \wedge A_1$ or $\forall x A_0$.

We say that A has $\bigwedge^+$-Form if it has $\bigwedge$-Form, is a true primitive recursive formula, or has the form $\mu xX.A(x,X)n$. Define

$$C[k] := \begin{cases} C_k & \text{if } C = C_0 \wedge C_1 \text{ or } C = C_0 \vee C_1 \text{ where } k \in \{0,1\} \\ A(k) & \text{if } C = \forall xA \text{ or } C = \exists xA \text{ where } k \in \mathbb{N} \end{cases}$$

Lemma 6.1. If C is a $\bigwedge$-Form then there is a uniform operator $\mathcal{J}^k_C$ such that whenever $d \vdash_m \Gamma, C$, $\mathcal{J}^k_C(d) \vdash_m \Gamma, C[k]$.

Proof. By induction on d.

$$\mathcal{J}^k_C(d) := \begin{cases} \mathcal{J}^k_C(d_k) & \text{if } \mathcal{I} = \bigwedge_C \\ CB_{F\mapsto\Sigma,C[k]}(\mathcal{J}^k_C(d_0)) & \text{if } \mathcal{I} = CB_{F\mapsto\Sigma,C} \\ \neg\{\mathcal{J}^k_C \circ \mathcal{F}_q\}_q & \text{if } \mathcal{I} = \neg\{\mathcal{F}_q\}_q \\ \mathcal{I}(\mathcal{J}^k_C(d_\iota))_{\iota\in|\mathcal{I}|} & \text{otherwise} \end{cases}$$

□

Lemma 6.2. Let $rk(C) \leq m$ with $\bigwedge^+$-Form and $e \vdash_m \Gamma, C$. Then there is an operator $\mathcal{R}_C(e, \cdot)$ such that whenever $d \vdash_m \Gamma, \neg C$, $\mathcal{R}_C(e, d) \vdash_m \Gamma$ and such that $\{\mathcal{R}_C\} \cup \{\mathcal{J}^k_D\}_{k,D}$ is uniform.

Proof. By induction on d.

$$\mathcal{R}_C(e,d) := \begin{cases} Cut_{C[k]}\mathcal{J}^k_C(e)\mathcal{R}_C(e,d_k) & \text{if } \mathcal{I} = \bigvee^k_{\neg C} \\ e & \text{if } \mathcal{I} = Ax_{\neg C, C} \\ CB_{F\mapsto\Sigma}\mathcal{R}_C(e,d_0) & \text{if } \mathcal{I} = CB_{F\mapsto\Sigma,\neg C} \\ \neg\{\mathcal{R}_C \circ \mathcal{F}_q\}_q & \text{if } \mathcal{I} = \neg\{\mathcal{F}_q\}_q \\ \mathcal{I}(\mathcal{R}_C(e,d_\iota))_{\iota\in|\mathcal{I}|} & \text{otherwise} \end{cases}$$

□

Lemma 6.3. For each m, there is an operator $\mathcal{E}_m$ so that whenever $d \vdash_{m+1} \Gamma$, $\mathcal{E}_m(d) \vdash_m \Gamma$ and $\{\mathcal{E}_m\} \cup \{\mathcal{R}_C\}_C \cup \{\mathcal{J}^k_D\}_{k,D}$ is uniform.

Proof. By induction on d.

$$\mathcal{E}_m(\mathcal{I}(d_\iota)_{\iota\in|\mathcal{I}|}) := \begin{cases} \mathcal{R}_C(\mathcal{E}_m(d_0), \mathcal{E}_m(d_1)) & \text{if } \mathcal{I} = Cut_C,\ rk(C) = m \\ & \text{ and } C \text{ has } \bigwedge^+\text{-Form} \\ \mathcal{R}_{\neg C}(\mathcal{E}_m(d_1), \mathcal{E}_m(d_0)) & \text{if } \mathcal{I} = Cut_C,\ rk(C) = m \\ & \text{ and } \neg C \text{ has } \bigwedge^+\text{-Form} \\ \neg\{\mathcal{E}_m \circ \mathcal{F}_q\}_q & \text{if } \mathcal{I} = \neg\{\mathcal{F}_q\}_q \\ \mathcal{I}(\mathcal{E}_m(d_\iota))_{\iota\in|\mathcal{I}|} & \text{otherwise} \end{cases}$$

□

Lemma 6.4. There is a uniform operator $\mathcal{D}_I$ such that if Γ does not contain predicates of inaccessible type negatively and $d \vdash_0 \Gamma$ then $\mathcal{D}_I(d) \vdash \Gamma$ and $\mathcal{D}_I(d) \in \mu^\infty_I$.

Proof. By induction on d.

$$\mathcal{D}_I(\mathcal{I}(d_\iota)_{\iota\in|\mathcal{I}|}) := \begin{cases} F \circ \mathcal{D}_I & \text{if } \mathcal{I} = \Omega_{n \notin \mu x X.A(x,X)} \\ & \text{and } \mu x X.A(x,X) \text{ has inaccessible type} \\ \neg\{\mathcal{D}_I \circ \mathcal{F}_q\}_q & \text{if } \mathcal{I} = \neg\{\mathcal{F}_q\}_q \\ \mathcal{I}(\mathcal{E}_m(d_\iota))_{\iota\in|\mathcal{I}|} & \text{otherwise} \end{cases}$$

□

Lemma 6.5. There is an operator $\mathcal{D}_n$ such that if $d \vdash_0 \Gamma$ and $dp(\Gamma) \leq n$ then $\mathcal{D}_n(d) \vdash_0 \Gamma$ and is a proof in ID_n^∞.

Proof. By induction on d.

$$\mathcal{D}_n(\mathcal{I}(d_\iota)_{\iota\in|\mathcal{I}|}) := \begin{cases} \mathcal{D}_n(d_{\mathcal{D}_m(d_0)}) & \text{if } \mathcal{I} = \Omega_{n \notin \mu x X.A(x,X)} \\ & \text{and } dp(\mu x X.A(x,X)) = m \geq n \\ \mathcal{D}_n(d_{\mathcal{D}_I(d_0)}) & \text{if } \mathcal{I} = \Omega_{n \in \mu x X.A(x,X)} \\ & \text{and } \mu x X.A(x,X) \text{ has inaccessible type} \\ \mathcal{I}(\mathcal{D}_n(d_\iota))_{\iota\in|\mathcal{I}|} & \text{otherwise} \end{cases}$$

□

Theorem 6.1. Let d be a proof in μ_2 of a sequent Γ of depth 0. Then there is a cut-free proof d^* of Γ in ID_0^∞. Furthermore, the existence may be shown in a constructive theory.

Proof. Let $d^* := \mathcal{D}_0(\mathcal{E}_0(\cdots(\mathcal{E}_m(d^\infty))))$. Then d^* is a cut-free proof in ID_0^∞.

Constructivity follows via continuous cut-elimination carried in an appropriate constructive system; for specificity, intuitionistic $\Pi^1_2 - CA$ would be (more than) sufficient to formalize each instance of this argument. Although the derivations are nominally infinite, they can be replaced with finitary descriptions, with branches only produced when they are actually used. Since all our transformations are defined continuously, they remain well-defined in this context. □

Theorem 6.2. μ_2 is consistent.

Proof. If there is a proof of $0 = 1$ in μ_2 then there is a cut-free proof in ID_0^∞. But the cut-free proofs of primitive recursive formulas are also proofs in IS, so there is a cut-free proof of $0 = 1$ in μ_2. But this is impossible, since no inference rule other than cut can produce this as an end-sequent. □

References

[1] K. Schütte, Beweistheoretische Erfassung der unerdlichen Induktion in der Zahlentheorie, Math. Ann. 122 (1951) 369–389.

[2] W. Buchholz, Eine Erweiterung der Schnitteliminationsmethode, Habilitationsschrift (1977).

[3] W. Pohlers, Cut elimination for impredicative infinitary systems. I. Ordinal analysis for ID_1, Arch. Math. Logik Grundlag. 21 (3-4) (1981) 113–129.

[4] M. Rathjen, Recent advances in ordinal analysis: Π^1_2-CA and related systems, Bull. Symbolic Logic 1 (4) (1995) 468–485.

[5] T. Arai, Proof theory for theories of ordinals. I. Recursively Mahlo ordinals, Ann. Pure Appl. Logic 122 (1-3) (2003) 1–85.

[6] T. Arai, Proof theory for theories of ordinals. II. Π_3-reflection, Ann. Pure Appl. Logic 129 (1-3) (2004) 39–92.

[7] T. Arai, Proof theory for theories of ordinals, preprints (1996-1997).

[8] G. Takeuti, M. Yasugi, The ordinals of the systems of second order arithmetic with the provably Δ^1_2-comprehension axiom and with the Δ^1_2-comprehension axiom respectively, Japan. J. Math. 41 (1973) 1–67.

[9] G. Jäger, W. Pohlers, Eine beweistheoretische Untersuchung von (Δ^1_2-CA)+(BI) und verwandter Systeme, Bayer. Akad. Wiss. Math.-Natur. Kl. Sitzungsber. (1982) 1–28 (1983).

[10] W. Pohlers, Subsystems of set theory and second order number theory, in: Handbook of proof theory, Vol. 137 of Stud. Logic Found. Math., North-Holland, Amsterdam, 1998, pp. 209–335.

[11] W. Buchholz, K. Schütte, Proof theory of impredicative subsystems of analysis, Vol. 2 of Studies in Proof Theory. Monographs, Bibliopolis, Naples, 1988.

[12] M. Möllerfeld, Generalized inductive definitions, Ph.D. thesis, Westfälische Wilhelms-Universität Münster (2003).

[13] W. Buchholz, Relating ordinals to proofs in a perspicuous way, in: Reflections on the foundations of mathematics (Stanford, CA, 1998), Vol. 15 of Lect. Notes Log., Assoc. Symbol. Logic, Urbana, IL, 2002, pp. 37–59.

[14] J. Barwise, L. S. Moss, Vicious Circles: On the Mathematics of Non-Wellfounded Phenomena, CSLI Publications, 1996.

[15] P. Aczel, Non-well-founded sets, Vol. 14 of CSLI Lecture Notes, Stanford University Center for the Study of Language and Information, Stanford, CA, 1988, with a foreword by Jon Barwise [K. Jon Barwise].

[16] A. Setzer, Well-ordering proofs for Martin-Löf type theory with W-type and one universe., Annals of Pure and applied Logic 92 (1998) 113–159.