On the Structure of Inductive Reasoning: Circular and Tree-Shaped Proofs in the μ-Calculus

Christoph Sprenger[1,*] and Mads Dam[2,**]

[1] INRIA Sophia Antipolis, France (sprenger@sophia.inria.fr)
[2] Royal Institute of Technology, Kista, Sweden (mfd@imit.kth.se)

Abstract In this paper we study induction in the context of the first-order μ-calculus with explicit approximations. We present and compare two Gentzen-style proof systems each using a different type of induction. The first is based on finite proof trees and uses a local well-founded induction rule, while the second is based on (finitely represented) ω-regular proof trees and uses a global induction discharge condition to ensure externally that all inductive reasoning is well-founded. We give effective procedures for the translation of proofs between the two systems, thus establishing their equivalence.

1 Introduction

Induction is the fundamental finitary method available in mathematics and computer science to generate and reason about finite and infinite objects. Three main proof-theoretic approaches to explicit induction[1] can be identified:

1. well-founded induction,
2. Scott/Park fixed point induction (cf. [7]), and
3. cyclic proofs, based on the idea of reducing induction to a global well-foundedness check on a locally sound inference graph.

As general approaches to inductive reasoning the former two are clearly the dominant ones. However, the value of cyclic reasoning in the decidable case has been demonstrated quite unequivocally by now. Examples are the well-established equivalence between monadic second-order logic, temporal logics (including various μ-calculi) and automata on infinite objects [14] as well as the usefulness of automata for obtaining decidability results and efficient algorithms [15]. Similar claims can be made concerning the usefulness of tableau-based techniques for

* Research done mainly while at Swedish Institute of Computer Science (SICS), supported by Swiss European Fellowship 83EU-065536, and completed at INRIA, supported by an ERCIM fellowship.

** Supported by the Swedish Research Council grant 621-2001-2637, "Semantics and Proof of Programming Languages"

[1] As opposed to implicit induction (cf. [3]), based on Knuth-Bendix completion.

obtaining completeness results in modal and temporal logic [6]. In the context of theorem proving and (undecidable) proof search, however, cyclic reasoning has received little attention. In our opinion, this situation deserves to be remedied. Our claim is that it facilitates proof search, mainly due to its ability to delay decisions concerning induction strategy as far as possible. Although it is too early for any real conclusions on the validity of this claim, the experiences with cyclic proofs for the μ-calculus using the EVT theorem prover [1,5] are encouraging.

In this paper, we address the fundamental question of the relative deductive power of cyclic proofs versus well-founded induction, the latter being the yardstick by which other formalisations must be compared. Our investigation is based on Park's first-order μ-calculus [8], which provides a minimal setting to study formalised induction. In this context, cyclic reasoning underpins work on model checking [13], completeness of the modal μ-calculus [16], and, more recently, Gentzen-type proof systems for compositional verification [4,11,10]. We establish effective translations between two Gentzen-style proof systems: one, $\mathcal{S}_{loc}$, for well-founded (local) induction, and the other, $\mathcal{S}_{glob}$, based on (finitely represented) ω-regular proof trees using an external global induction discharge condition ensuring the well-foundedness. We work in an extension of the basic μ-calculus with explicit approximations [4] and ordinal quantification [10] (Sect. 2). Inductive reasoning in both proof systems rests on this extension. In system $\mathcal{S}_{loc}$ (Sect. 3), it is supported by a single local induction rule, an instantiation of the the well-known rule of well-founded induction to ordinals. In system $\mathcal{S}_{glob}$ (Sect. 4), the global induction discharge condition organises the basic cycles, called *repeats*, into a partial *induction order,* assigns a progressing induction variable to each repeat and requires each repeat to preserve (i.e. not increase) the variables of repeats above it in the induction order. This condition ensures that the induction associated with each strongly connected subgraph is well-founded. For the translation from $\mathcal{S}_{loc}$ to $\mathcal{S}_{glob}$ (Sect. 5) it is sufficient to show that the local induction rule of $\mathcal{S}_{loc}$ is derivable in $\mathcal{S}_{glob}$. The translation in the other direction (Sect. 6) is more involved and generally proceeds in two stages. We first present a direct translation for $\mathcal{S}_{glob}$-proofs, where the inductive structure represented in the induction order matches the proof tree structure. Then, we discuss an exponential time algorithm, which unfolds arbitrary cyclic proofs until they are in the form required by the direct translation.

We think that, by clearly exhibiting the underlying structures and their relationships, our present formal treatment sheds some new light on the various approaches to inductive reasoning. An important benefit from using explicit approximations is that it largely decouples our constructions from the actual language (here, the μ-calculus), thus strongly suggesting that they can support lazy-style global induction in other contexts such as type theories with inductive definitions [9]. Barthe et al. [2], for instance, points in this direction by proposing a type system ensuring termination of recursive functions based on approximations of inductive types. Finally, an interesting practical implication of our result is that (assuming size blow-ups can be prevented) it permits standard theorem provers to be gracefully interfaced with the μ-calculus based EVT prover.

Set-theoretic preliminaries Let $G = (A, R \subseteq A \times A)$ be a graph. We say that G is a *forest* if $(a, b) \in R$ and $(a, c) \in R$ imply $b = c$. A *tree* is a forest with a unique *root node* $r \in A$ such that there is no $a \in A$ with $(a, r) \in R$. We call G *forest-like* if $(a, b) \in R$ and $(a, c) \in R$ imply $b = c$ or $(b, c) \in R \cup R^{-1}$. A subset $C \subseteq A$ is *strongly R-connected* if for any two $x, y \in C$ we have that $(x, y) \in R^*$. C is *weakly R-connected* if $R \cup R^{-1}$ is strongly R-connected. We sometimes call $C \subseteq A$ a subgraph of G and mean the subgraph $(C, R \cap C \times C)$ induced by C. A strongly connected subgraph (SCS) $C \subseteq A$ is *non-trivial* if $R \cap C \times C \neq \varnothing$.

2 μ-Calculus with Explicit Approximations

Let $\mathbf{2} = \{0, 1\}$ be the two-point lattice and let $\mathsf{Pred}(S) = \mathbf{2}^S$ be the lattice of predicates over S ordered pointwise. For a monotone map $f \colon \mathsf{Pred}(S) \to \mathsf{Pred}(S)$ we define the *ordinal approximation* $\mu^\alpha f$ and the *fixed point* μf by

$$\begin{array}{ll} \mu^0 f \quad = \lambda x.0 & \mu^\gamma f = \bigvee_{\alpha<\gamma} \mu^\alpha f \quad \text{for limit ordinals } \gamma \\ \mu^{\alpha+1} f = f(\mu^\alpha f) & \mu f \quad = \bigvee_a \mu^\alpha f \end{array}$$

Proposition 1. *Let $f \colon \mathsf{Pred}(S) \to \mathsf{Pred}(S)$ be a monotone map. Then*

1. *μf is the least fixed point of f (Knaster-Tarski), and*
2. *$\mu^\alpha f = \bigvee_{\beta<\alpha} f(\mu^\beta f)$.*

We assume countably infinite sets of individual variables $x, y, z \ldots \in V_I$, of predicate variables $X, Y, Z, \ldots \in V_P$ of each arity $n \geq 0$, and of ordinal variables $\iota, \kappa, \lambda, \ldots \in V_O$. Let t range over the terms of some signature Σ.

Definition 1. (Syntax) *The syntax of* μ-calculus formulas ϕ and predicates Φ over Σ *is inductively defined by*

$$\begin{array}{l} \phi ::= t = t' \mid \kappa' < \kappa \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \exists x.\phi \mid \exists\kappa.\phi \mid \Phi(\overline{t}) \\ \Phi ::= X \mid \mu X(\overline{x}).\phi \mid \mu^\kappa X(\overline{x}).\phi \end{array}$$

with the restriction that the arities of Φ and $\overline{t}$ match in $\Phi(\overline{t})$ and the formation of both $\mu X(\overline{x}).\phi$ and $\mu^\kappa X(\overline{x}).\phi$ is subject to the conditions that (i) the arities of X and $\overline{x}$ match, and (ii) all occurrences of X in ϕ appear under an even number of negations (formal monotonicity).

Zero-ary predicates are identified with formulas. We call formulas of the form $\kappa' < \kappa$ *ordinal constraints*. The sets of free and bound variables of formulas and predicates are defined as expected, in particular, $\mathrm{fv}(\mu X(\overline{x}).\phi) = \mathrm{fv}(\phi) - \{X, \overline{x}\}$ and $\mathrm{fv}(\mu X^\kappa(\overline{x}).\phi) = (\mathrm{fv}(\phi) - \{X, \overline{x}\}) \cup \{\kappa\}$. We will identify formulas that are equal up to renaming of bound variables. Dual connectives are defined in the usual way, with $\nu X(\overline{x}).\phi = \neg\mu X(\overline{x}).\neg\phi[\neg X/X]$ and the approximated fixed point similarly. We also use bounded ordinal quantifiers defined by the abbreviations $\exists\iota<\kappa.\phi = \exists\iota.\, \iota<\kappa \wedge \phi$ and $\exists\iota\leq\kappa.\phi = (\exists\iota<\kappa.\phi) \vee \phi[\kappa/\iota]$.

Given a Σ-model $\mathcal{M} = (\mathcal{A}, \rho)$ ($\mathcal{A}$ is the structure and ρ the interpretation) the semantics interprets a μ-calculus formula ϕ as an element $\|\phi\|_\mathcal{M} \in \mathbf{2}$ and an n-ary predicate Φ as an element $\|\Phi\|_\mathcal{M} \in \mathsf{Pred}(|A|^n)$. We often drop $\mathcal{M}$ and write $\|\phi\|_\rho$ and $\|\Phi\|_\rho$ if the structure $\mathcal{A}$ is clear from the context.

Definition 2. (Semantics) *Given a signature Σ and a Σ-model $(\mathcal{A}, \rho)$ define the semantics of μ-calculus formulas ϕ and predicates Φ over Σ inductively by*

$$\begin{array}{llll}
\|t = t'\|_\rho & = \mathit{if}\ \|t\|_\rho = \|t'\|_\rho\ \mathit{then}\ 1\ \mathit{else}\ 0 & \|\exists\kappa.\phi\|_\rho & = \bigvee_\beta \|\phi\|_{\rho[\beta/\kappa]} \\
\|\kappa' < \kappa\|_\rho & = \mathit{if}\ \rho(\kappa') < \rho(\kappa)\ \mathit{then}\ 1\ \mathit{else}\ 0 & \|\Phi(\overline{t})\|_\rho & = \|\Phi\|_\rho(\|\overline{t}\|_\rho) \\
\|\neg\phi\|_\rho & = 1 - \|\phi\|_\rho & \|X\|_\rho & = \rho(X) \\
\|\phi_1 \wedge \phi_2\|_\rho & = \min\{\|\phi_1\|_\rho, \|\phi_2\|_\rho\} & \|\mu X(\overline{x}).\phi\|_\rho & = \mu\Psi \\
\|\exists x.\phi\|_\rho & = \bigvee_{a \in |A|} \|\phi\|_{\rho[a/x]} & \|\mu^\kappa X(\overline{x}).\phi\|_\rho & = \mu^{\rho(\kappa)}\Psi
\end{array}$$

where $\|\overline{t}\|_\rho$ is defined as usual and $\Psi = \lambda P.\lambda\overline{a}.\|\phi\|_{\rho[P/X,\overline{a}/\overline{x}]}$ in the clauses for fixed point and approximation predicates.

Given a model $\mathcal{M} = (\mathcal{A}, \rho)$, we extend the valuation ρ a posteriori to terms t and formulas ϕ by defining $\rho(t) = \|t\|_\rho$ and $\rho(\phi) = \|\phi\|_\rho$. This allows us to compose substitutions θ with environments ρ as in $\rho \circ \theta$. We say that a model $\mathcal{M} = (\mathcal{A}, \rho)$ *satisfies* a formula ϕ, written $\mathcal{M} \models \phi$, if $\|\phi\|_\rho = 1$. A formula ϕ is *valid* if $\mathcal{M} \models \phi$ for all models $\mathcal{M}$.

3 Local Induction: the System $\mathcal{S}_{loc}$

In this section we introduce the Gentzen-type proof system $\mathcal{S}_{loc}$ for local well-founded induction. It shares most definitions and proof rules with the system $\mathcal{S}_{glob}$ for global induction presented in the next section.

Sequents The *sequents* of both proof systems are of the form $\Gamma \vdash_\mathcal{O} \Delta$, where Γ and Δ are finite multisets of formulas and $\mathcal{O}$ is a finite set of ordinal variables. A sequent is *well-formed* if all ordinal variables occurring free in Γ or Δ are elements of $\mathcal{O}$. We tacitly restrict our attention to well-formed sequents. The set of free variables of a sequent is defined by $\mathrm{fv}(\Gamma \vdash_\mathcal{O} \Delta) = \mathrm{fv}(\Gamma \cup \Delta) \cup \mathcal{O}$. Substitutions are extended to multisets of formulas by defining $\Gamma[\theta] = \{\phi[\theta] \mid \phi \in \Gamma\}$. In sequents we often write $\mathcal{O}, \kappa$ for $\mathcal{O} \cup \{\kappa\}$.

Given a Σ-model $\mathcal{M} = (\mathcal{A}, \rho)$ we say that $\mathcal{M}$ *satisfies* a sequent $\Gamma \vdash_\mathcal{O} \Delta$ if $\mathcal{M} \models \phi$ for all $\phi \in \Gamma$ implies that $\mathcal{M} \models \psi$ for some $\psi \in \Delta$. A model $\mathcal{M}$ *falsifies* a sequent $\Gamma \vdash_\mathcal{O} \Delta$ if $\mathcal{M}$ does not satisfy $\Gamma \vdash_\mathcal{O} \Delta$. The sequent $\Gamma \vdash_\mathcal{O} \Delta$ is *valid* if it is satisfied in all models.

Proof System The proof system is presented in two parts. Table 1 shows the basic set of proof rules, common to both $\mathcal{S}_{loc}$ and $\mathcal{S}_{glob}$. They are presented in tableau-style with the conclusion above the line and the premises below. Fixed point rule (μ_1-L) is the essential device which introduces ordinal approximations. Note the asymmetry of rules (μ_1-L) and (μ_0-R). However, one can show that the symmetric rules (μ_0-L) and (μ_1-R) are derivable in $\mathcal{S}_{loc}$. The local proof system $\mathcal{S}_{loc}$ is then obtained by adding to the basic proof rules the local induction rule (Ind-L) of Table 2. The table also shows its derivable dual (Ind-R), which might look more familiar to the reader.

Structural Rules

$$(\text{Id})\ \frac{\Gamma, \phi \vdash_{\mathcal{O}} \phi, \Delta}{\cdot} \qquad (\text{Weak})\ \frac{\Gamma \vdash_{\mathcal{O}} \Delta}{\Gamma' \vdash_{\mathcal{O}'} \Delta'}\ \begin{array}{l}\Gamma' \subseteq \Gamma \\ \Delta' \subseteq \Delta \\ \mathcal{O}' \subseteq \mathcal{O}\end{array}$$

$$(\text{Cut})\ \frac{\Gamma \vdash_{\mathcal{O}} \Delta}{\Gamma, \phi \vdash_{\mathcal{O}} \Delta \quad \Gamma \vdash_{\mathcal{O}} \phi, \Delta} \qquad (\text{Subst})\ \frac{\Gamma[\theta] \vdash_{\mathcal{O}[\theta]} \Delta[\theta]}{\Gamma \vdash_{\mathcal{O}} \Delta}$$

Logical and Equality Rules

$$(\neg\text{-L})\ \frac{\Gamma, \neg\phi \vdash_{\mathcal{O}} \Delta}{\Gamma \vdash_{\mathcal{O}} \phi, \Delta} \qquad (\neg\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} \neg\phi, \Delta}{\Gamma, \phi \vdash_{\mathcal{O}} \Delta}$$

$$(\wedge\text{-L})\ \frac{\Gamma, \phi_1 \wedge \phi_2 \vdash_{\mathcal{O}} \Delta}{\Gamma, \phi_1, \phi_2 \vdash_{\mathcal{O}} \Delta} \qquad (\wedge\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} \phi_1 \wedge \phi_2, \Delta}{\Gamma \vdash_{\mathcal{O}} \phi_1, \Delta \quad \Gamma \vdash_{\mathcal{O}} \phi_2, \Delta}$$

$$(\exists\text{-L})\ \frac{\Gamma, \exists x.\phi \vdash_{\mathcal{O}} \Delta}{\Gamma, \phi \vdash_{\mathcal{O}} \Delta}\ x \notin \text{fv}(\Gamma, \Delta) \qquad (\exists\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} \exists x.\phi, \Delta}{\Gamma \vdash_{\mathcal{O}} \phi[t/x], \Delta}$$

$$(=\text{-L})\ \frac{\Gamma[t_2/x], t_1 = t_2 \vdash_{\mathcal{O}} \Delta[t_2/x]}{\Gamma[t_1/x] \vdash_{\mathcal{O}} \Delta[t_1/x]} \qquad (=\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} t = t, \Delta}{\cdot}$$

Fixed Point Rules

$$(\mu_1\text{-L})\ \frac{\Gamma, (\mu X(\overline{x}).\phi)(\overline{t}) \vdash_{\mathcal{O}} \Delta}{\Gamma, \exists\kappa.(\mu^{\kappa} X(\overline{x}).\phi)(\overline{t}) \vdash_{\mathcal{O}} \Delta} \qquad (\mu_0\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} (\mu X(\overline{x}).\phi)(\overline{t}), \Delta}{\Gamma \vdash_{\mathcal{O}} \phi[\mu X(\overline{x}).\phi/X, \overline{t}/\overline{x}], \Delta}$$

$$(\mu^{\kappa}\text{-L})\ \frac{\Gamma, (\mu^{\kappa} X(\overline{x}).\phi)(\overline{t}) \vdash_{\mathcal{O}} \Delta}{\Gamma, \exists\kappa' < \kappa.\phi[\mu^{\kappa'} X(\overline{x}).\phi/X, \overline{t}/\overline{x}] \vdash_{\mathcal{O}} \Delta}$$

$$(\mu^{\kappa}\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} (\mu^{\kappa} X(\overline{x}).\phi)(\overline{t}), \Delta}{\Gamma \vdash_{\mathcal{O}} \exists\kappa' < \kappa.\phi[\mu^{\kappa'} X(\overline{x}).\phi/X, \overline{t}/\overline{x}], \Delta}$$

Ordinal Rules

$$(\exists\kappa\text{-L})\ \frac{\Gamma, \exists\kappa.\phi \vdash_{\mathcal{O}} \Delta}{\Gamma, \phi \vdash_{\mathcal{O},\kappa} \Delta}\ \kappa \notin \mathcal{O} \qquad (\exists\kappa\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} \exists\kappa.\phi, \Delta}{\Gamma \vdash_{\mathcal{O}} \phi[\iota/\kappa], \Delta}\ \iota \in \mathcal{O}$$

$$(<\text{-L})\ \frac{\Gamma, \kappa < \kappa \vdash_{\mathcal{O}} \Delta}{\cdot} \qquad (<\text{-R})\ \frac{\Gamma \vdash_{\mathcal{O}} \kappa' < \kappa, \Delta}{\Gamma \vdash_{\mathcal{O}} \kappa' < \kappa'', \Delta \quad \Gamma \vdash_{\mathcal{O}} \kappa'' < \kappa, \Delta}$$

Table 1. The proof rules shared by $\mathcal{S}_{loc}$ and $\mathcal{S}_{glob}$

$$\text{(Ind-L)}\ \frac{\Gamma, \exists\kappa.\phi \vdash_{\mathcal{O}} \Delta}{\Gamma, \phi \vdash_{\mathcal{O},\kappa} \exists\kappa' < \kappa.\phi[\kappa'/\kappa], \Delta}\ \kappa \notin \mathcal{O}$$

$$\text{(Ind-R)}\ \frac{\Gamma \vdash_{\mathcal{O}} \forall\kappa.\phi, \Delta}{\Gamma, \forall\kappa' < \kappa.\phi[\kappa'/\kappa] \vdash_{\mathcal{O},\kappa} \phi, \Delta}\ \kappa \notin \mathcal{O}$$

Table 2. The local induction rules of $\mathcal{S}_{loc}$

Given a set $\mathcal{S}$ of proof rules an $\mathcal{S}$-*derivation tree* $\mathcal{D} = (\mathcal{N}, \mathcal{E}, \mathcal{L})$ is a tree $(\mathcal{N}, \mathcal{E})$ with nodes $\mathcal{N}$ and edges $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ together with a function $\mathcal{L}$ labelling each node of the tree with a sequent in a way that is consistent with the application of the proof rules in $\mathcal{S}$. We often write $N(\Gamma \vdash_{\mathcal{O}} \Delta)$ for $\mathcal{L}(N) = \Gamma \vdash_{\mathcal{O}} \Delta$.

Definition 3. *An* $\mathcal{S}_{loc}$-proof for a sequent $\Gamma \vdash_{\mathcal{O}} \Delta$ *is an* $\mathcal{S}_{loc}$*-derivation tree* $\mathcal{D}$ *whose root is labelled by* $\Gamma \vdash_{\mathcal{O}} \Delta$ *and each leaf of which is labelled by an axiom.*

Lemma 1. *The proof rules of Tables 1 and 2 are sound. In particular, if there is a* Σ*-model* $(\mathcal{A}, \rho)$ *falsifying the conclusion* C *of a rule then there is a environment* ρ' *such that* $(\mathcal{A}, \rho')$ *falsifies some premise* P *of that rule. Moreover, we can choose* ρ' *such that, for all ordinal variables* κ *free in both* P *and* C*, we have* $\rho'(\kappa) = \rho(\theta(\kappa))$ *for rule (Subst) and* $\rho'(\kappa) \leq \rho(\kappa)$ *for all other rules.*

The proof proceeds by a straightforward inspection of the rules. In particular, the soundness of the fixed point rules follows immediately from Proposition 1. The soundness of the proof system $\mathcal{S}_{loc}$ is then an immediate corollary.

Theorem 1. (Soundness of $\mathcal{S}_{loc}$) *If* $\mathcal{D}$ *is an* $\mathcal{S}_{loc}$*-proof for* $\Gamma \vdash_{\mathcal{O}} \Delta$ *then* $\Gamma \vdash_{\mathcal{O}} \Delta$ *is valid.*

4 Global Induction: the System $\mathcal{S}_{glob}$

The proof system $\mathcal{S}_{glob}$ uses the proof rules from Table 1 only. However, proofs in this system are not finite but ω-regular trees, which we represent as finite trees with back edges called *repeats*. An external global *induction discharge condition* then ensures that all inductive reasoning embodied in the proof structure is well-founded. Let us fix an arbitrary $\mathcal{S}_{glob}$-derivation tree $\mathcal{D} = (\mathcal{N}, \mathcal{E}, \mathcal{L})$.

Definition 4. (Repeat) *A* repeat $R = (N, M)$ *for* $\mathcal{D}$ *is a pair of nodes of* $\mathcal{D}$ *such that* N *is a leaf,* M *lies on the path from the root of* $\mathcal{D}$ *to* N *and* $\mathcal{L}(N) = \mathcal{L}(M)$*. The node* N *is called the* repeat node *and* M *is called its* companion*. We denote by* $\pi(R)$ *the path* $M \cdots N$ *in* $\mathcal{D}$*.*

Definition 5. (Pre-Proof) *A* pre-proof $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ *is composed of an* $\mathcal{S}_{glob}$*-derivation tree* $\mathcal{D} = (\mathcal{N}, \mathcal{E}, \mathcal{L})$ *and a set of repeats* $\mathcal{R}$ *for* $\mathcal{D}$ *such that each of its non-axiom leaves appears in exactly one repeat of* $\mathcal{R}$*. We call the graph* $\mathcal{G}(\mathcal{P}) = (\mathcal{N}, \mathcal{E} \cup \mathcal{R}, \mathcal{L})$ *the* pre-proof graph *of* $\mathcal{P}$*.*

The following lemma will allow us to identify strongly connected subgraphs of $\mathcal{G}(\mathcal{P})$ with certain subsets of $\mathcal{R}$. For two repeats $R = (N, M)$ and $R' = (N', M')$ in $\mathcal{R}$ define $R \to R'$ if there is a path $M \cdots N'$ from the companion node M of R to the repeat node N' of R' in the derivation tree $\mathcal{D}$.

Lemma 2. *There is a bijection between the non-trivial strongly connected subgraphs of $\mathcal{G}(\mathcal{P})$ and the strongly connected subgraphs of $(\mathcal{R}, \to)$.*

We are now ready to define the *basic induction discharge condition* qualifying a pre-proof as a proof. This condition is based on the notions of preservation and progress of repeats with respect to approximant variables.

Definition 6. (Progress, Preservation) *Constraint $\kappa' < \kappa$ is called* derivable *at $N(\Gamma \vdash_{\mathcal{O}} \Delta)$, written $N \vdash \kappa' < \kappa$, if there is a repeat-free $\mathcal{S}_{glob}$-pre-proof for $\Gamma \vdash_{\mathcal{O}} \kappa' < \kappa, \Delta$. Let R be a repeat with $\pi(R) = N_0 \cdots N_m$ and $N_i(\Gamma_i \vdash_{\mathcal{O}_i} \Delta_i)$ and let κ be an ordinal variable. Then we say*

– R preserves κ, *if $\kappa \in \mathcal{O}_i$ for all i, and if either $N_j \vdash \theta(\kappa) < \kappa$ or $\theta(\kappa) = \kappa$ whenever rule (Subst) is applied with θ at N_j, and*
– R progresses on κ, *if R preserves κ and rule (Subst) is applied with some θ at some N_j such that $N_j \vdash \theta(\kappa) < \kappa$.*

Definition 7. ($\mathcal{S}_{glob}$-Proof) *A pre-proof $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ is an $\mathcal{S}_{glob}$-*proof *if for each strongly connected subgraph $S \subseteq \mathcal{R}$ there is an ordinal variable κ such that*

1. *some repeat $R \in S$ progresses on κ, and*
2. *each repeat $R' \in S$ preserves κ.*

Theorem 2. (Soundness of $\mathcal{S}_{glob}$) *If $\mathcal{P}$ is an $\mathcal{S}_{glob}$-proof for $\Gamma \vdash_{\mathcal{O}} \Delta$ then $\Gamma \vdash_{\mathcal{O}} \Delta$ is valid.*

Proof. (Sketch) By contradiction. Using Lemma 1 we construct an infinite sequence $(N_0, \rho_0) \cdots (N_i, \rho_i) \cdots$ of pairs of nodes and valuations such that each ρ_i falsifies $\mathcal{L}(N_i)$. By the definition of $\mathcal{S}_{glob}$-proof there is an ordinal variable κ such that the sequence $\{\rho_i(\kappa)\}_i$ of ordinals decreases infinitely often from some point on, contradicting the well-foundedness of the ordinals. □

Discharge Using Induction Orders The basic induction discharge mechanism does not exhibit sufficient structure for our purpose of comparing the two systems $\mathcal{S}_{loc}$ and $\mathcal{S}_{glob}$. For this reason we introduce an alternative induction discharge condition, first proposed by Schöpp [10] and generalising the one in [4], which orders the set of repeats appearing in a pre-proof. The new condition turns out to be equivalent to the original one.

Definition 8. (Induction Orders) *A partial order $(\mathcal{R}, \preceq)$ on the set of repeats is called an* induction order *for $\mathcal{P}$, if it is forest-like and every strongly connected subgraph $S \subseteq \mathcal{R}$ has a $\preceq$-greatest element.*

A labelled induction order *$(\mathcal{R}, \preceq, \delta)$ is an induction order $(\mathcal{R}, \preceq)$ together with a map δ assigning an ordinal variable κ to each repeat $R \in \mathcal{R}$. The ordinal variable $\delta(R)$, also written δ_R, is called the* induction variable *for R.*

The restriction to forest-like partial orders in this definition is adopted here for convenience. It is not necessary for soundness, but sufficient for completeness (see Proposition 2 below).

Definition 9. (Alternative Discharge) *We say that a labelled induction order* $(\mathcal{R}, \preceq, \delta)$ discharges *a pre-proof* $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ *if for all* $R \in \mathcal{R}$

1. *R progresses on δ_R, and*
2. *R preserves $\delta_{R'}$ whenever $R \preceq R'$.*

Proposition 2. *For any pre-proof $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ the following are equivalent:*

(i) there is a labelled induction order $(\mathcal{R}, \preceq, \delta)$ discharging $\mathcal{P}$, and
(ii) $\mathcal{P}$ is an $\mathcal{S}_{glob}$-proof.

5 From Local to Global Induction

The translation of $\mathcal{S}_{loc}$-proofs to $\mathcal{S}_{glob}$-proofs is achieved by showing that the local induction rule (Ind-L) is derivable in $\mathcal{S}_{glob}$, in the strong sense that any application of (Ind-L) inside an $\mathcal{S}_{glob}$-proof can be replaced by an equivalent $\mathcal{S}_{glob}$-derivation.

Theorem 3. *The local induction rule* (Ind-L) *is derivable in* $\mathcal{S}_{glob}$.

Proof. Consider the following derivation (omitting two applications of the weakening rule):

$$\dfrac{\dfrac{\Gamma, \exists\kappa.\phi \vdash_{\mathcal{O}} \Delta}{\lceil \Gamma, \phi \vdash_{\mathcal{O},\kappa} \Delta \rceil}\,(\exists\kappa\text{-L})}{\Gamma, \phi \vdash_{\mathcal{O},\kappa} \exists\kappa' < \kappa.\phi[\kappa'/\kappa], \Delta \qquad \dfrac{\dfrac{\Gamma, \exists\kappa' < \kappa.\phi[\kappa'/\kappa] \vdash_{\mathcal{O},\kappa} \Delta}{\Gamma, \kappa' < \kappa, \phi[\kappa'/\kappa] \vdash_{\mathcal{O},\kappa,\kappa'} \Delta}\,(\exists\kappa\text{-L},\wedge\text{-L})}{\lfloor \Gamma, \phi \vdash_{\mathcal{O},\kappa} \Delta \rfloor}\,(\text{Subst})}\,(\text{Cut})$$

This derivation is sound provided $\kappa \notin \mathcal{O}$, the side condition of rule (Ind-L). Since the repeat (indicated by brackets) preserves all variables in $\mathcal{O}$ and progresses on κ, this derivation can safely replace applications of (Ind-L) in $\mathcal{S}_{glob}$-proofs. □

6 From Global to Local Induction

In general, the translation from $\mathcal{S}_{glob}$-proofs to $\mathcal{S}_{loc}$-proofs proceeds in two stages. If the inductive structure of the $\mathcal{S}_{glob}$-proof matches its tree structure, it can be translated directly into an $\mathcal{S}_{loc}$-proof. Otherwise, the $\mathcal{S}_{glob}$-proof needs to be unfolded prior to this transformation. We fix an arbitrary pre-proof $\mathcal{P} = (\mathcal{D}, \mathcal{R})$.

Definition 10. (Structural dependency) *The* structural dependency relation $\leq_{\mathcal{P}}$ *on $\mathcal{R}$ is defined as follows: $R' \leq_{\mathcal{P}} R$ holds for two repeats $R, R' \in \mathcal{R}$ if the companion node of R' appears on the path $\pi(R)$.*

Lemma 3. *Let $S \subseteq \mathcal{R}$. Then S is strongly $\rightarrow$-connected if and only if S is weakly $\leq_{\mathcal{P}}$-connected.*

Definition 11. (Tree-compatibility) *An induction order $(\mathcal{R}, \preceq)$ for $\mathcal{P}$ is* tree-compatible *if $R \leq_{\mathcal{P}} R'$ and $R' \not\leq_{\mathcal{P}} R$ imply $R \preceq R'$ for all $R, R' \in \mathcal{R}$. An $\mathcal{S}_{glob}$-proof $\mathcal{P}$ is called* tree-dischargeable *if there is a tree-compatible induction order discharging $\mathcal{P}$.*

The inductive structure of a tree-dischargeable proof matches its underlying tree structure. The next lemma, which can be proved using Lemma 3, gives a useful characterisation of induction orders for $\mathcal{P}$ in terms of the structural dependency relation $\leq_{\mathcal{P}}$, thereby relating the structure of the proof tree with the dependencies between repeats in an arbitrary induction order.

Lemma 4. *Let $(\mathcal{R}, \preceq)$ be a forest-like partial order. Then $(\mathcal{R}, \preceq)$ is an induction order for $\mathcal{P}$ if and only if $R \leq_{\mathcal{P}} R'$ implies $R \preceq R'$ or $R' \preceq R$ for all $R, R' \in \mathcal{R}$.*

Let $\preceq_{\mathcal{P}}$ be the transitive closure of $\leq_{\mathcal{P}}$. The following two remarks are easy corollaries of Lemma 4.

Proposition 3. *Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a pre-proof such that $\mathcal{R}$ is* injective *(as a function from repeat nodes to companions). Then $\preceq_{\mathcal{P}}$ is a tree-compatible induction order for $\mathcal{P}$.*

Lemma 5. *Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a proof discharged by the tree-compatible induction order $(\mathcal{R}, \preceq, \delta)$. Then $\mathcal{P}$ can be transformed into a proof $\mathcal{P}' = (\mathcal{D}', \mathcal{R}')$ of the same sequent such that $\mathcal{R}'$ is injective and $\mathcal{P}'$ is discharged by $(\mathcal{R}', \preceq_{\mathcal{P}'}, \delta')$ for some labelling δ'.*

6.1 Translating Tree-Dischargeable Proofs

Since each repeat R embodies an induction progressing on variable δ_R along the path $\pi(R)$ from the companion to the repeat node, it seems natural to insert the local induction rule (Ind-L) at the companion node and use, essentially, the whole sequent as an induction hypothesis. This induction hypothesis is then conveyed down the proof tree, exploiting progress on δ_R along $\pi(R)$ to remove the bounded quantification introduced by (Ind-L) and thus making the induction hypothesis available for local discharge at the repeat node. The rest of this section is devoted to proving the following theorem along these lines.

Theorem 4. *Let $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ be a $\mathcal{S}_{glob}$-proof of $\Gamma \vdash_{\mathcal{O}} \Delta$, tree-dischargeable by $(\mathcal{R}, \preceq, \delta)$. Then $\mathcal{P}$ can be transformed into a $\mathcal{S}_{loc}$-proof $\mathcal{D}'$ of $\Gamma \vdash_{\mathcal{O}} \Delta$.*

Proof. (Sketch) We assume w.l.o.g. that (A) $\mathcal{R}$ is injective and $\preceq = \preceq_{\mathcal{P}}$, by Lemma 5, and (B) companions appear only as descendents of nodes where a rule other than (Subst), (Cut), ($\wedge$-R) and (<-R) is applied.

Our construction recursively transforms $\mathcal{D} = (\mathcal{N}, \mathcal{E}, \mathcal{L})$ into a new derivation tree $\mathcal{D}' = (\mathcal{N}', \mathcal{E}', \mathcal{L}')$ by replacing at each node N the rule applied to produce the set of descendents $\mathcal{N}_N$ by a derivation $\mathcal{D}_N$ with root $\widehat{N}$ and premises $\widehat{\mathcal{N}_N} = \{\widehat{N'} \mid N' \in \mathcal{N}_N\}$ (and some fresh interior nodes). The procedure keeps the set $\mathcal{H}$ of current *induction hypotheses*, which is added to $N(\Gamma \vdash_{\mathcal{O}} \Delta)$ yielding $\widehat{N}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H})$ in $\mathcal{D}'$. For any node $N \in \mathcal{N}$ define the set of repeats *active* at N by $\mathcal{R}_N = \{R' \in \mathcal{R} \mid \exists R \in \mathcal{R}. N \in \pi(R) \text{ and } R \preceq_{\mathcal{P}} R'\}$ and let $\mathcal{O}_N = \{\delta_R \mid R \in \mathcal{R}_N\}$ be the corresponding set of induction variables. We note fact (C): there is some $R \in \mathcal{R}_N$ preserving all variables in $\mathcal{O}_N$. Hence, $\mathcal{O}_N \subseteq \mathcal{O}$ for $N(\Gamma \vdash_{\mathcal{O}} \Delta)$.

The induction hypothesis H_R to be added to the current set of hypotheses $\mathcal{H}$ at the companion node $M(\Gamma \vdash_{\mathcal{O}} \Delta)$ of a repeat $R = (N, M)$ is defined by

$$\begin{aligned} H_R &= \exists \delta < \delta_R . \Psi_R[\delta/\delta_R] \\ \Psi_R &= \exists \bar{\iota} \leq \bar{\iota}_R . \exists \bar{v} . \bigwedge(\Gamma \cup \neg\Delta \cup \neg\mathcal{H})[\bar{\iota}/\bar{\iota}_R] \end{aligned}$$

where $\{\bar{\iota}_R\} = \mathcal{O}_M - \{\delta_R\}$ and $\{\bar{v}\} = \mathrm{fv}(\Gamma \vdash_{\mathcal{O}} \Delta) - \mathcal{O}_M$. The *free* induction hypothesis Ψ_R packs the sequent at M together with the set $\mathcal{H}$ into a single formula, existentially quantifies all but the active induction variables in $\mathcal{O}_M$ and binds the (preserved) active induction variables $\bar{\iota}_R$ in Ψ_R by a quantifier of type $\exists \cdot \leq \bar{\iota}_R$. The *guarded* induction hypothesis H_R additionally binds the (progressing) induction variable δ_R in Ψ_R by a quantifier of type $\exists \cdot \leq \delta_R$.

Transformation of $\mathcal{D}$ *to* $\mathcal{D}'$. Our procedure ensures that for each node $N(\Gamma \vdash_{\mathcal{O}} \Delta)$ in $\mathcal{D}$ there is a node $\widehat{N}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H})$ such that the following invariant holds:

$$I(N, \mathcal{H}) = \ \mathcal{H}_N^g \subseteq \mathcal{H} \subseteq \mathcal{H}_N$$

where $\mathcal{H}_N^g = \{H_R \mid R \in \mathcal{R}_N\}$ and $\mathcal{H}_N = \mathcal{H}_N^g \cup \{\Psi_R \mid R \in \mathcal{R}_N\}$. By assumption (B) the root N_r is not a companion so $\mathcal{R}_N = \varnothing$ and the invariant holds trivially by initially setting $\mathcal{H} = \varnothing$. We now describe the derivations $\mathcal{D}_N$ replacing in $\mathcal{D}'$ the rule application at N in $\mathcal{D}$. Suppose we have constructed $\mathcal{D}'$ up to some node $\widehat{N}(\Gamma' \vdash_{\mathcal{O}'} \Delta', \mathcal{H})$ where the invariant holds. The cases where a rule other than (Subst) is applied at N and none of the descendents of N is a companion are easy to show. We just remark that in order to maintain the invariant at the branching rules (Cut), ($\wedge$-R) and (<-R) we possibly need weakening on $\mathcal{H}$ to account for the splitting of the set of active repeats between the two branches. Due to assumption (B) the two remaining cases are:

Case 1. The single descendent $M(\Gamma \vdash_{\mathcal{O}} \Delta)$ *of* N *is the companion of a repeat* R. The invariant $I(M, \mathcal{H})$ is violated for $\mathcal{H}$, because the induction hypothesis H_R is missing in $\mathcal{H}$. The derivation $\mathcal{D}_N$ in Fig. 1 adds H_R to $\widehat{N}(\Gamma' \vdash_{\mathcal{O}'} \Delta', \mathcal{H})$ yielding $\widehat{M}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}, H_R)$ and thus reestablishing the invariant.

At node M' we cut in $\exists \delta_R . \Psi_R$. After weakening away all but the latter formula on the left hand side, we apply the induction rule (Ind-L). This leaves us with the sequent $\Psi_R \vdash_{\mathcal{O}_N} H_R$, which is transformed into the desired sequent $\widehat{M}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}, H_R)$ by applying a series of essentially first-order rules (RS1) deconstructing Ψ_R on the left. On the right hand side, we apply a dual series of rules (RS2) proving $N_r(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}, \Psi_R)$ locally. Note that the sequent at N_r anticipates, thanks to the cut, the desired situation at the repeat node of R, since it contains a right hand side occurrence of the free induction hypothesis.

$$\cfrac{\cfrac{\widehat{N}(\Gamma' \vdash_{\mathcal{O}'} \Delta', \mathcal{H})}{M'(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H})}\ \text{(rule at } N\text{)}}{\cfrac{\cfrac{\cfrac{\Gamma, \exists \delta_R.\Psi_R \vdash_{\mathcal{O}} \Delta, \mathcal{H}}{\exists \delta_R.\Psi_R \vdash_{\mathcal{O}_N - \{\delta_R\}}}\ \text{(Weak)}}{\Psi_R \vdash_{\mathcal{O}_N} H_R}\ \text{(Ind-L)}}{\widehat{M}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}, H_R)}\ \text{(RS1)} \qquad \cfrac{\cfrac{\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}, \exists \delta_R.\Psi_R}{N_r(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}, \Psi_R)}\ (\exists\kappa\text{-R})}{\cdot}\ \text{(RS2)}}\ \text{(Cut)}$$

Figure 1. Derivation $\mathcal{D}_N$ inserted at companion node N of a repeat R

Case 2. Rule (Subst) applied at N with descendent M. We need to make sure that the substitution rule is correctly applied and that the induction hypotheses in $\mathcal{H}$ are (re-)generated as in the following compressed version of derivation $\mathcal{D}_N$:

$$\cfrac{\cfrac{\widehat{N}(\Gamma[\theta] \vdash_{\mathcal{O}[\theta]} \Delta[\theta], \mathcal{H})}{M'(\Gamma[\theta] \vdash_{\mathcal{O}[\theta]} \Delta[\theta], \mathcal{H}'[\theta])}\ \text{(Regen)}}{\widehat{M}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}')}\ \text{(Subst)}$$

where $\mathcal{H}' = \{H_R \mid H_R \in \mathcal{H}\} \cup \{\Psi_R \mid \Psi_R \in \mathcal{H} \text{ or } N \vdash \theta(\delta_R) < \delta_R\}$. We have $\mathcal{R}_M = \mathcal{R}_N$, since M is not a companion by assumption (B). It is then easy to see that $I(M, \mathcal{H}')$ holds. Note that, since $\mathrm{fv}(\mathcal{H}_N) \subseteq \mathcal{O}_N$, it follows from fact (C) above that each $\kappa \in \mathrm{fv}(\mathcal{H}_N)$ is preserved by θ at N, that is, $N \vdash \theta(\kappa) < \kappa$ or $\theta(\kappa) = \kappa$. The derivation from $\widehat{N}$ to M' labelled (Regen) then includes, for each $R \in \mathcal{R}_N$, a derivation that produces:

1. $\Psi_R[\theta]$ from H_R if $N \vdash \theta(\delta_R) < \delta_R$, exploiting progress of δ_R and preservation of $\overline{\iota}_R$ by θ at N,
2. $\Psi_R[\theta]$ from Ψ_R if $\theta(\delta_R) = \delta_R$ and $\Psi_R \in \mathcal{H}$, using preservation of δ_R and $\overline{\iota}_R$ by θ at N, and
3. $H_R[\theta]$ from H_R, also using preservation of δ_R and $\overline{\iota}_R$ by θ at N.

In each of the derivations (1)-(3) the leading bounded ordinal quantifiers in Ψ_R and H_R are duplicated and commuted prior to instantiation as necessary, by applying some easily derivable auxiliary rules.

Continuing this procedure down to the leaves of $\mathcal{D}$ yields, for each repeat $R = (N, M)$, two nodes $\widehat{M}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H})$ and $\widehat{N}(\Gamma \vdash_{\mathcal{O}} \Delta, \mathcal{H}')$ in $\mathcal{D}'$ (where $\widehat{N}$ is a leaf of $\mathcal{D}'$ so far). We now show that $\mathcal{H} \subseteq \mathcal{H}'$ and $\Psi_R \in \mathcal{H}'$, implying that the sequent at $\widehat{N}$ is provable in $\mathcal{S}_{loc}$ in the same way as the one at node N_r in Fig. 1. Consider some $R' \in \mathcal{R}_M$. From $\mathcal{R}_M = \mathcal{R}_N$ and the invariant it follows that $H_{R'}$ is in both $\mathcal{H}_M$ and $\mathcal{H}_N$. If $\Psi_{R'} \in \mathcal{H}_M$ then we also have $\Psi_{R'} \in \mathcal{H}_N$, since R preserves $\delta_{R'}$ and so $\Psi_{R'}$ is regenerated along $\pi(R)$ (see discussion in case (2) above). Hence, $\mathcal{H} \subseteq \mathcal{H}'$. Since R progresses on δ_R, Ψ_R will be generated from H_R at some point on the path $\pi(R)$ and then regenerated in each subsequent application of rule (Subst) along $\pi(R)$. Hence, we have $\Psi_R \in \mathcal{H}'$. This shows that $\mathcal{D}'$ can be completed into an $\mathcal{S}_{loc}$-proof. $\square$

6.2 General Case: Unfolding Proofs

The previous translation crucially depends on the tree-dischargeability of the induction order: repeats with companions lower in the proof tree preserve induction variables of repeats higher in the proof tree ("higher" and "lower" being determined by $\preceq_{\mathcal{P}}$). In general, we need to unfold the proof until it becomes tree-dischargeable. This task is achieved by Algorithm 1.

```
input
  $\mathcal{P}_0 = (\mathcal{D}_0, \mathcal{R}_0)$ where $\mathcal{D}_0 = (\mathcal{N}_0, \mathcal{E}_0, \mathcal{L}_0)$, root $N_r$  { $\mathcal{R}_0$ injective }
  $(\mathcal{R}_0, \preceq_0, \delta_0)$  { induction order discharging $\mathcal{P}_0$ }
output
  $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ where $\mathcal{D} = (\mathcal{N}, \mathcal{E}, \mathcal{L})$, $\mathcal{N} \subseteq \mathcal{N}_0 \times \mathbb{N}$  { unfolded proof }
globals
  $s \in \mathbb{N}$  { sequence number to distinguish copies of nodes }
begin
  $s := 0$; $\mathcal{E} := \varnothing$; $\mathcal{R} := \varnothing$
  $\mathcal{N} := \{(N_r, s)\}$; $\mathcal{L} := \{((N_r, s), \mathcal{L}_0(N_r))\}$
  unfold $(N_r, s)$ $\varnothing$
end

procedure unfold $(N, k)$ $\mathcal{B}$
parameters
  $(N, k) \in \mathcal{N}_0 \times \mathbb{N}$  { node of $\mathcal{P}$ = (node of $\mathcal{P}_0$, copy number) }
  $\mathcal{B} \colon \mathcal{R}_0 \rightharpoonup \mathbb{N}$  { copy numbers for companions available in $\mathcal{P}$ }
if $N$ is the repeat node of some $R = (N, M) \in \mathcal{R}_0$ then
  if $(R, i) \in \mathcal{B}$ for some $i$ then { companion $(M, i)$ available for $(N, k)$ }
    $\mathcal{R} := \mathcal{R} \cup \{((N, k), (M, i))\}$
  else { no companion available, continue unfolding through repeat }
    $s := s + 1$  { get a new sequence number }
    add node $(M, s)$ labelled $\mathcal{L}_0(M)$ as child of $(N, k)$ to $\mathcal{D}$
    unfold $(M, s)$ $\mathcal{B}$
  end if
else { $N$ is an axiom leaf or an inner node of $\mathcal{D}_0$ }
  if $N$ is the companion of some $R \in \mathcal{R}_0$ and $R \notin \operatorname{dom} \mathcal{B}$ then
    $\mathcal{B} := \{(R, k)\} \cup \{(R', k') \in \mathcal{B} \mid R \preceq_0 R'\}$
  end if
  for each child $N'$ of $N$ in $\mathcal{D}_0$ do { add and unfold each child node }
    add node $(N', k)$ labelled $\mathcal{L}_0(N')$ as child of $(N, k)$ to $\mathcal{D}$
    unfold $(N', k)$ $\mathcal{B}$
  end for
end if
```

Algorithm 1: Unfolding proofs

It takes a proof $\mathcal{P}_0 = (\mathcal{D}_0, \mathcal{R}_0)$ with injective $\mathcal{R}_0$ as input and produces a tree-dischargeable proof $\mathcal{P} = (\mathcal{D}, \mathcal{R})$. Note that no generality is lost by requiring that $\mathcal{R}_0$ is injective. The nodes of $\mathcal{P}$ are pairs (N, k), where N is a node of $\mathcal{P}_0$ and k is the copy number. The original proof tree is traversed recursively, unfolding repeats into new copies of the proof tree as necessary. The procedure maintains a partial map $\mathcal{B}$ from repeats $\mathcal{R}_0$ to copy numbers in $\mathbb{N}$ to keep track of companions that are available for looping back at repeat nodes. This map is updated whenever we encounter the companion of some repeat $R \in \mathcal{R}$ without an entry in $\mathcal{B}$: the entry (R, k) is added to $\mathcal{B}$, while any entry for a repeat R' not above R with respect to $\preceq_0$ is removed from $\mathcal{B}$ (line 26-27). When examining copy l of the repeat node N of some repeat $R = (N, M) \in \mathcal{R}$ we check whether there is some entry $(R, k) \in \mathcal{B}$ (lines 17-18). If so, then we can safely close the loop and add $((N, l), (M, k))$ as a new repeat to $\mathcal{R}$ (line 19). Otherwise, if there is no companion available for R, we proceed by unfolding the tree at (N, l) by adding the node (M, k) with a fresh k as a descendant of (N, l) to $\mathcal{P}$ (line 21-23).

The labelled induction order $(\mathcal{R}, \preceq, \delta)$ for $\mathcal{P}$ is obtained by lifting $(\mathcal{R}_0, \preceq_0, \delta_0)$ to $\mathcal{P}$. Writing $\widehat{R} = (N, M) \in \mathcal{R}_0$ for $R = ((N, k), (M, i)) \in \mathcal{R}$, we define

$$R \preceq R' \Leftrightarrow \widehat{R} \prec_0 \widehat{R'} \text{ or } (\widehat{R} = \widehat{R'} \text{ and } k \leq k') \quad \text{and} \quad \delta(R) = \delta_0(\widehat{R})$$

where $R = ((N, k), (M, i))$ and $R' = ((N', k'), (M', i'))$. Note the tie-breaking role of the repeat sequence number in case of identical projected repeats.

Theorem 5. *Let $\mathcal{P}_0 = (\mathcal{D}_0, \mathcal{R}_0)$ be a $\mathcal{S}_{glob}$-proof of $\Gamma \vdash_{\mathcal{O}} \Delta$ such that $\mathcal{R}_0$ is injective and $\mathcal{P}_0$ is discharged by $(\mathcal{R}_0, \preceq_0, \delta_0)$. Then Algorithm 1 produces, in time $\mathcal{O}(2^{|\mathcal{N}_0| \times |\mathcal{R}_0|})$, a proof $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ of $\Gamma \vdash_{\mathcal{O}} \Delta$, discharged by the tree-compatible induction order $(\mathcal{R}, \preceq, \delta)$ defined above.*

Proof. For partial correctness it is sufficient to show that

1. $\mathcal{P} = (\mathcal{D}, \mathcal{R})$ is a pre-proof,
2. $(\mathcal{R}, \preceq)$ is a tree-compatible induction order for $\mathcal{P}$, and
3. $(\mathcal{R}, \preceq, \delta)$ discharges $\mathcal{P}$.

It is easy to see that $\mathcal{P}$ is a pre-proof. As a preparation for (2) and (3) consider a repeat $R = ((N, k), (M, i))$ in $\mathcal{R}$. Let

$$c = ((N_0, k_0), \mathcal{B}_0) \cdots ((N_j, k_j), \mathcal{B}_j) \cdots ((N_m, k_m), \mathcal{B}_m)$$

with $(N_0, k_0) = (M, i)$ and $(N_m, k_m) = (N, k)$ be the (suffix of the) sequence of recursive calls leading to the introduction of R in $\mathcal{P}$. The pair $(\widehat{R}, i)$ is added to $\mathcal{B}_0$ in call c_0 (line 27) and appears in all subsequent $\mathcal{B}_j$ $(1 \leq j \leq m)$. Finally, the repeat R is added to $\mathcal{R}$ in call c_m (line 19). Consider a repeat $R'' = (N'', M'') \in \mathcal{R}_0$ with its companion $M'' = N_j$ occurring in c_j with $0 \leq j \leq m$. Since $(\widehat{R}, i)$ is in both $\mathcal{B}_j$ and $\mathcal{B}_{j+1}$ and (R'', k_j) is added to $\mathcal{B}_j$ at in call c_j (line 27), we have

(A) if $R'' \notin \operatorname{dom} \mathcal{B}_j$ then $R'' \preceq_0 \widehat{R}$.

Ad (2). It is easy to check that $(\mathcal{R}, \preceq, \delta)$ is a forest-like partial order. Now suppose $R' \leq_{\mathcal{P}} R$ for some $R, R' \in \mathcal{R}$. Since the companion of R' appears on $\pi(R)$ in $\mathcal{P}$ and hence as N_j in some c_j it must be the case that $\widehat{R'} \notin \operatorname{dom} \mathcal{B}_j$ so we have $\widehat{R'} \preceq_0 \widehat{R}$ by (A). From the definition of $\preceq$ and Lemma 4 it follows that $(\mathcal{R}, \preceq)$ is an induction order. If, moreover, $R \not\leq_{\mathcal{P}} R'$ then $\widehat{R'} \neq \widehat{R}$, since $\widehat{R}$ cannot be unfolded on $\pi(R)$. This implies that $(\mathcal{R}, \preceq)$ is tree-compatible.

Ad (3). Let $S' \subseteq \mathcal{R}_0$ be the set of repeats unfolded on $\pi(R)$ and let $S = S' \cup \{\widehat{R}\}$. Note that $\pi(R)$ in $\mathcal{P}$ is the composition of all the paths $\pi(R'')$ in $\mathcal{P}_0$ with $R'' \in S$. Suppose $R \preceq R'$ for some $R' \in \mathcal{R}$. Then certainly $\widehat{R} \preceq_0 \widehat{R'}$. Let $R'' \in S'$ with companion M. Since R'' is unfolded in some call c_{j-1}, we have $N_j = M$ and $R'' \notin \operatorname{dom} \mathcal{B}_j$. Hence, $R'' \preceq_0 \widehat{R}$ by (A) and R'' preserves $\delta(R') = \delta_0(\widehat{R'})$, implying that R preserves $\delta(R')$. Moreover, R progresses on $\delta(R)$, since $\widehat{R}$ progresses on $\delta_0(\widehat{R})$. This shows (3).

Complexity. Suppose for a contradiction that there is a (suffix of a) sequence of calls of the form c above such that $m > 0$ and $(N_0, \operatorname{dom} \mathcal{B}_0) = (N_m, \operatorname{dom} \mathcal{B}_m)$. Note first that, since the control flow of `unfold` does not depend on the copy numbers, there is an extension $((N_{m+1}, k_{m+1}), \mathcal{B}_{m+1}) \cdots ((N_{2m}, k_{2m}), \mathcal{B}_{2m})$ of c such that $(N_i, \operatorname{dom} \mathcal{B}_i) = (N_{i+m}, \operatorname{dom} \mathcal{B}_{i+m})$ for all $i \leq m$ (and, in fact, so on ad infinitum). Let S be the set of repeats $R \in \mathcal{R}_0$ such that N_i is the companion of R and $R \notin \operatorname{dom} \mathcal{B}_i$ for some $i \leq m$. It is not difficult to see that S is strongly connected in $(\mathcal{R}_0, \rightarrow)$ and thus there is a $\preceq_0$-greatest element $\widetilde{R} \in S$. Suppose N_i is the companion of $\widetilde{R}$ for some $i \leq m$. Since $\widetilde{R} \notin \operatorname{dom} \mathcal{B}_i$ and $\widetilde{R}$ is $\preceq_0$-greatest in S, it follows that $\widetilde{R} \in \operatorname{dom} \mathcal{B}_k$ for all $i < k \leq i + m$ (line 27). In particular, $\widetilde{R} \in \operatorname{dom} \mathcal{B}_{i+m}$, which contradicts $\operatorname{dom} B_i = \operatorname{dom} B_{i+m}$. Hence, length of any call sequence c of `unfold` is bounded by $|c| \leq |\mathcal{N}_0| \times |\mathcal{R}_0|$, yielding an upper bound of $|\mathcal{N}| \leq 2^{|\mathcal{N}_0| \times |\mathcal{R}_0|}$ for the size of $\mathcal{P}$ and the time complexity of the algorithm. $\square$

7 Conclusions

We have presented a translation between proofs using well-founded induction and cyclic proofs based on a global well-foundedness condition. The proof systems use explicit ordinal approximations as suggested in [4]. Since our main interest in approximants is as a proof-theoretical mechanism to deal with fixed points rather than proving theorems about them, it would be desirable to identify a fragment of the language which could be shown to conserve (not increase) the expressiveness of the basic μ-calculus (without explicit approximations). Simpson and Schöpp have proposed an alternative approach to approximants based directly on the second-order variables instead of ordinal variables [11] and they have proved a conservativity result for a variant of this language [12]. Their language lacks, however, the ability to "internalise" sequents into single formulas, required in our direct translation to local proofs. Thus, it seems that our approach can not readily be adapted to yield a similar translation in their framework.

On a different line of research, we would like to investigate whether the ideas of this paper can be transferred to the context of type theories with inductive definitions such as the Calculus of Inductive Constructions [9]. A useful starting point is the introduction of approximated inductive types along the lines of [2].

Acknowledgements We would like to thank Alex Simpson and Uli Schöpp as well as the members of the FDT group at SICS for fruitful discussions on the topic. We are also grateful to Dilian Gurov, Marieke Huisman and the anonymous referees for their helpful suggestions.

References

1. T. Arts, G. Chuganov, M. Dam, L.-å. Fredlund, D. Gurov, and T. Noll. A tool for verifying software written in Erlang. Accepted for publication in *STTT Journal*, 2001.
2. G. Barthe, M. J. Frade, E. Giménez, L. Pinto, and T. Uustalu. Type-based termination of recursive definitions. *Mathematical Structures in Computer Science*, 2000. to appear.
3. H. Comon. Inductionless induction. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 14. Elsevier Science, 2001.
4. M. Dam and D. Gurov. μ-calculus with explicit points and approximations. *Journal of Logic and Computation*, 12(2):43–57, 2002. Previously appeared in Fixed Points in Computer Science, FICS '02.
5. L. Fredlund. *A Framework for Reasoning about Erlang Code*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2001.
6. R. Goré. Tableau methods for modal and temporal logics. In *Handbook of Tableau Methods*. Kluwer, 1999.
7. D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.
8. D. Park. Finiteness is mu-ineffable. *Theoretical Computer Science*, 3(2):173–181, 1976.
9. C. Paulin-Mohring. Inductive definitions in the system Coq – rules and properties. Technical Report 92-49, Laboratoire de l'Informatique du Parallélisme, ENS Lyon, France, Dec. 1992.
10. U. Schöpp. Formal verification of processes. Master's thesis, University of Edinburgh, 2001.
11. U. Schöpp and A. Simpson. Verifying temporal properties using explicit approximants: Completeness for context-free processes. In *FOSSACS '02, Grenoble, France*, volume 2303 of *LNCS*, pages 372–386. Springer-Verlag, 2002.
12. A. Simpson and U. Schöpp. Private communication.
13. C. Stirling and D. Walker. Local model checking in the modal μ-calculus. *Theoretical Computer Science*, 89:161–177, 1991.
14. W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 133–191. Elsevier Science Publishers, Amsterdam, 1990.
15. M. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Logic in Computer Science, LICS '86*, pages 322–331, 1986.
16. I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. In *Logic in Computer Science, LICS '95*, pages 14–24, 1995.