Felix Joachimski

Reduction Properties of ΠIE-Systems

Herbert Utz Verlag · Wissenschaft
München

Dissertation zur Erlangung des Grades eines
Doktors der Naturwissenschaften
an der Fakultät für Mathematik und Informatik
der Ludwig-Maximilians-Universität München

Eingereicht am 13. 11. 2001

1. Berichterstatter: Prof. Dr. Helmut Schwichtenberg
2. Berichterstatter: Prof. Dr. Wilfried Buchholz

Tag der mündlichen Prüfung: 21.12.2001

Die Deutsche Bibliothek – CIP-Einheitsaufnahme

Ein Titeldatensatz für diese Publikation ist
bei Der Deutschen Bibliothek erhältlich

Zugleich: Dissertation, München, Univ., 2001

ISBN 3-8316-0105-4

Printed in Germany

Herbert Utz Verlag GmbH, München
Tel.: 089/277791-00 – Fax: 089/277791-01

Contents

Introduction

This thesis presents so-called ΠIE-systems and their coinductive analogues and studies their operational behavior from a basic proof-theoretical perspective. The first part is mainly concerned with the untyped phenomena of confluence, standardization and inductive characterizations of terms in general ΠIE-systems. The second part focuses on the normalization properties of some well-known example calculi and in particular their typed versions.

Term systems with binders

ΠIE-systems form a subset of the set of higher order infinitary term systems. These are distinguished from first order infinitary term systems by their mechanism of variable binding that internalizes substitution in the shape of abstraction.

Finite term systems with binders are commonly implemented in the meta-language of the λ-calculus by means of higher-order abstract syntax [PE88]. The framework of term systems introduced in subsection 2.1 allows for infinitely branching terms and thus would not fit into finitely presentable λ-calculi. They should rather be classified as semi-formal systems (in the sense of proof theory [Sch77]), enriched and complicated by the presence of variable binders that require a careful handling of variables and substitution for them. In particular, the various variable conventions (see [Hin97] for a thorough discussion), frequently employed to efficiently avoid variable name clashes for finite calculi, fail in infinitary systems. As a consequence we have to resort to a de Bruijn-style variable management [Bru72] with its inherently complex notions of lifting, substitution etc. The theory of de Bruijn-systems would probably find its most accurate description in terms of modern rank 2 inductive definitions ("abstract syntax" [FPT99, AR99, BP99]), but we have chosen the more conventional rank 1 approach, since its exposition is just as simple and requires less involved meta-theoretic means.

ΠIE-systems

The acronym ΠIE abridges "permutative introduction/elimination", alluding to the fundamental dichotomy in natural deduction between constructors, which *introduce* concepts, and destructors (*eliminations*), which define their use. Indeed, most λ-calculi that correspond to natural deduction derivations via the Curry-Howard-translation, are instances of ΠIE-systems.

A natural deduction elimination rule consists of one main premise and some additional side premises, typically of the form

$$\dfrac{\begin{array}{c} \\ \vdots \\ A \end{array} \quad \begin{array}{c} B_1 \\ \vdots \\ C_1 \end{array} \quad \begin{array}{c} B_2 \\ \vdots \\ C_2 \end{array}}{D}$$

β-reduction. If the main premise is itself obtained by an introduction rule, this may be simplified [Pra65] by means of a computational (also called β)

reduction rule. The form of the result is, of course, dictated by the semantical goal of correctness, but nevertheless can be outlined in syntactical terms, so that essential operational properties such as confluence are ensured.

In subsection 3.3 we provide the relevant terminology to achieve such an abstract notion of β-conversion, thereby refining Ruckert's conception of β-reduction for the general natural deduction systems of his thesis [Ruc85]. In fact, the very notion of $\Pi\mathsf{IE}$-systems has been stimulated by his reflections, and our proceeding uses his groundwork as a starting point.

Permutations. Consider an elimination scheme with the same conclusion as (some) of its side premises which becomes the main premise of another elimination:

$$\dfrac{\dfrac{\begin{matrix}\vdots\\ A\end{matrix} \quad \begin{matrix}B_1\\ \vdots\\ C_1\end{matrix} \quad \begin{matrix}B_2\\ \vdots\\ C_2\end{matrix}}{C_1} \quad \begin{matrix}B_1'\\ \vdots\\ C_1'\end{matrix} \quad \begin{matrix}B_2'\\ \vdots\\ C_2'\end{matrix}}{D}$$

It is often reasonable to move the second premise into the side premise of the first elimination. In the above example derivation this results in

$$\dfrac{\begin{matrix}\vdots\\ A\end{matrix} \quad \dfrac{\begin{matrix}B_1\\ \vdots\\ C_1\end{matrix} \quad \begin{matrix}B_1'\\ \vdots\\ C_1'\end{matrix} \quad \begin{matrix}B_2'\\ \vdots\\ C_2'\end{matrix}}{D} \quad \begin{matrix}B_2\\ \vdots\\ C_2\end{matrix}}{D}$$

Such permutations have been introduced for disjunctive connectives like $+$ and $\exists$, but also make sense and are necessary in the normalization of proofs with the infinitely branching ω-rule [Maa75].

Permutations seem operationally harmless; their interaction with β-reduction is not very complex either. Surprisingly, it turns out that this intuition is rather difficult to substantiate, so that the mathematically most demanding argument of this thesis attends to a confluence proof for permutations in $\Pi\mathsf{IE}$-systems.

Non-wellfounded terms

Coinductive structures provide a natural environment for the semantics of infinite objects such as streams, or runs of an automaton. In this thesis, we apply some of the methods established in coalgebra [Rut00] to the structure of $\Pi\mathsf{IE}^{\mathrm{co}}$-systems that arise through a coinductive interpretation of the grammar of $\Pi\mathsf{IE}$-systems. These calculi of possibly non-wellfounded terms give rise to interesting phenomena; for instance, the coinductive analogue, Λ^{co}, of the λ-calculus Λ admits a direct definition of the recursion operator[2] $Yr := r(Yr)$ and embeds Böhm trees.

[2] This definition of Y — in contrast to the coded versions — is valid also in a typed system with coinductive type assignment rules and finite types.

The appropriate notions of substitution, conversion and reduction for such calculi are not immediate and require a careful balance of coinductive and inductive components. Particularly important is the question of how to adjust the concept of reduction sequence: while [KKSdV97] and the theory of infinitary term rewriting use infinitely many sequential reductions (and thus have to waive confluence), we choose infinitely many parallel conversions for each of only a finite number of reduction steps. The resulting theory can be developed parallelly to the wellfounded case and retains the fundamental theorems.

Bounds

As will be argued in subsection 2.3.1, the set of non-wellfounded terms can be recovered as the topological closure of its wellfounded subset. Accordingly, our extension of confluence and standardization results to non-wellfounded $\Pi\mathsf{IE}^{\mathrm{co}}$-systems may be interpreted as a proof of continuity for the underlying functions. Indeed, the main methodological achievements are bounding mechanisms that yield a modulus of continuity for the constructive content of the respective theorems. A fortiori, the bounds so established hold just as well for the wellfounded calculi, so that the pertinent sections can be read as a detailed analysis of the conventional proof methods from the viewpoint of complexity.

Notation systems for reduction relations

In customary expositions of confluence and standardization proofs, the underlying operations on reduction sequences are hidden inside the various (and often omitted) cases of several lemmas. The manipulation of reduction sequences is usually described in terms of geometric intuitions that lead to figures with many dots and lines and become quite bothersome upon formalization.[3]

For our task of adapting the confluence proofs to the non-wellfounded setting and extracting bounds, it is imperative that all computations on reduction (sequences) be explicit and demonstrably guarded recursive. To this end, a linear notation system for reduction, developments, standard reduction and the like is introduced as one of the main technical achievements, operating as a major tool throughout the first part of the thesis.

Confluence

The abstract confluence proof is divided into three parts: β-reductions, permutations and their combination.

In our treatment of **β-reductions** we generalize the traditional Tait/Martin-Löf development argument to the class of bounded $\Pi\mathsf{IE}$-systems, which are characterized by a uniform bound on the complexity of their various β-reducts ("substitution forms"). Takahashi's beautiful method [Tak95, MP93] — which uses complete developments — would be preferable, but it fails in infinitary

[3] Unfortunately, most of the actually formalized and implemented confluence proofs (see, e.g., [Nip01], where also many references are listed) are not concerned with program extraction.

ΠIE-systems, where the complete development is not necessarily attained by finite reduction sequences.

The confluence proof for **permutations** is a non-trivial generalization of a method first presented in [JM00] for the calculus ΛJ. That method heavily relied on the fact that permutations normalize and therefore the permutative normal form passes as a complete development for permutative reduction sequences. Again, such normal forms exist in ΠIE-systems[4] with sufficient commutation conditions, but may not be reached. As a remedy, they are replaced by approximations, stratified along the natural numbers for the case of ΛJ, and along bounded trees in the general case.

For the **combined reduction relation** we can once more avail ourselves of an idea of [JM00], where commutation of β-development and permutative reduction sequences served to connect the two confluence results. The relevant restrictions on the interaction between the two reduction relations are quite liberal for general ΠIE-systems, thereby complicating the commutation proof. Yet the overall structure of the confluence argument is preserved.

Standardization

Traditionally, the standardization theorem [Bar84] for the λ-calculus has been expressed in terms of residuals of reduction sequences. In [JM00], Ralph Matthes and the author gave an inductive characterization of the underlying notion of standardness that allows to shorten the proof and adapts to permutative conversions as well as to other situations [Ves01]. In this thesis' section on standardization the claim of [JM00] that the method translates to more complex calculi is substantiated. In fact, by a bounding argument in a spirit similar to that of the confluence section, it even extends to the non-wellfounded case, so that the length of the resulting standard reduction sequences can be estimated.[5]

Inductive characterizations

ΠIE-systems unite the features of many different calculi. Consequently, their basic definitions are inevitably complex and call for sufficiently powerful syntactic abstractions in order to make them at least presentable and perhaps even digestable. As a consequence, notational conventions and overloading abound in this thesis, serving to hide the underlying diversity in simple notation. More than a technical necessity, it is this meta-mathematical machinery that makes the theorems perspicuous.

One example figures prominently in the inductive (or wellfounded) part of this thesis: the uniform depiction of iterated eliminations R in form of the vector notation $r\vec{R}$. Its use has been advocated before in [JM] and continues to rule in this thesis, leading to tight inductive characterizations for terms and their normal forms along the head redex structure. These in turn can be

[4] They do not exist in ΠIE^{co}-systems.

[5] For the case of the λ-calculus, [Xi99] provided a complexity analysis, although neither method nor results are comparable with the considerations in this thesis.

extended to similar characterizations, WN, for the set of weakly and, SN, for that of the strongly normalizing terms. The former has been exemplified for the calculus $\Lambda\mathsf{J}$ (see below) in [JM00], where a completeness proof by help of standardization has been developed which now serves as a starting point for the non-trivial generalization to $\mathsf{\Pi IE}$-systems. Several instances of the latter were presented in [JM], but the abstract scheme is first formulated in the section on normalization.

Strong normalization

Most of the second part of the thesis is dedicated to variations on a proof technique for strong normalization that has repeatedly appeared in the literature [San67, Dil68, How80, Sch93] throughout the past four decades and enjoys renewed interest in recent work [Bec01, BW00]. In [JM] it has been shown that the use of SN considerably shortens the underlying argument and is necessary in order to conceive further applications, e.g., to permutative conversions. In the treatment of the λ-calculus (section 9) we retrace the essential steps and reformulate them, until bounds for the height of the reduction tree can be extracted. While the conventional proof-theoretic approach to such bounds uses the height of derivation trees,[6] it is likewise possible — albeit slightly surprising — to employ the size of derivations and this furnishes even sharper bounds.

η-expansion

η-expansion has been proposed by [Pra71] and studied by Mints in a Russian article of 1979, which passed unnoticed until the type theory community developed renewed interest for the expansive orientation of η-equality in the last decade. Various proofs of confluence and normalization for simply typed λ-calculi have been published (references will be provided in section 9) and also higher type systems have been explored. As the latest development, Barthe suggested a method for proving strong normalization for mixed β-reduction and η-expansion in the λ-cube [Bar99b]. Yet, essential proofs are omitted or only sketched, and they are very hard to reconstruct.

In the section on the λ-calculus, we recount the history and dangers of η-expansion and provide two different proofs of strong normalization. From an aesthetic point of view, the first approach is more appealing — it reveals that all η-expansions can be performed before any β-step is executed. However, this result does not extend to higher type systems like system F or the Calculus of Constructions. The second approach more subtly analyzes the interaction between β-reduction and η-expansion, utilizing the fact that η-expansion is a subset of the converse of η-reduction without overlapping with inverse β. The resulting commutation properties are formulated w.r.t. β-reduction and converse η-contraction only, and therefore apply to the untyped λ-calculus, from which they carry over to higher type systems and permit a purely syntactic and perspicuous strong normalization proof for the combined reduction relation in β-normalizing Pure Type Systems, where η-expansion is only slightly restricted.

[6] For Λ this has been demonstrated in [Bec01] and [Sch91].

The λ-calculus

The λ-calculus is the heart of this thesis.

First, it constitutes the major example, to which every method we discuss has to tune. Second, the theorems to be proved are essentially extensions of the theory of the λ-calculus. Third, the techniques exerted are derived from basic constructions in the λ-calculus.

This propinquity of our venture with the very elaborate theory of so classic a calculus is not to be considered a drawback, but an aid in keeping track of our position in the abstract theory.

More specifically, the contributions to the theory of the λ-calculus are the following:

- We prove bounds for the lengths of joining (parallel) reduction sequences in the confluence section and similar bounds for the lengths of standard reduction sequences in the standardization part.
- The coinductive (sometimes also called infinitary) calculus Λ^{co} is shown to be confluent and standardizing.
- Inductive characterizations for terms, normal forms, weakly and strongly normalizing terms are reproved.
- The normalization arguments of [JM] are repeated and carefully analyzed with regard to new bounds for the height of the reduction tree.
- Commutation of η-expansion with β-reduction is studied and strong normalization of their combination established by a new proof method.

The calculus $\Lambda\mathsf{J}$

Just as the λ-calculus prototypes β-reduction, the system $\Lambda\mathsf{J}$ serves as a minimal environment for the untyped study of permutations (which until the introduction of $\Lambda\mathsf{J}$ existed only in a typed world). All techniques for permutations are tested against $\Lambda\mathsf{J}$ before they can be applied to the general permutations of $\Pi\mathsf{IE}$-systems. As a consequence, many ideas for $\Pi\mathsf{IE}$-systems are derived from the two fundamental papers [JM00, JM], in which Ralph Matthes and the author outlined some aspects of the meta-theory of $\Lambda\mathsf{J}$.[7]

Confluence and standardization for $\Lambda\mathsf{J}$ and its coinductive analogue $\Lambda\mathsf{J}^{\mathrm{co}}$ follow from the results of the respective sections on $\Pi\mathsf{IE}^{\mathrm{co}}$-systems. In order to illustrate applications of a general permutation lemma for the set SN, we also added a short note on the strong normalization proof that follows [JM].

System T

We use the λ-based analogue of Gödel's system T [Göd58] as an example to demonstrate both the versatility of the concept of $\Pi\mathsf{IE}$-systems and the

[7] Further results have been obtained in [Mat00] and recently in [Mat01], where interpolation for the typed version of $\Lambda\mathsf{J}$ is validated.

strength of our proof method for strong normalization. The respective section massages the argument given in [JM] until a version is obtained that allows easy assignment of ordinals. Unfortunately, this final step that would lead to ordinal bounds for the complexity of functions that compute the reduction height of terms, is not carried out, although it would be particularly interesting to see whether the derivation size might again be preferable over the height, just as it has been for Λ.

Outline of the contents

Section 1 is a prerequisite for the whole thesis and should be read carefully, because it contains the central notions of sequences together with many notational conventions. Also, some terminology for coinduction and guarded recursion is fixed.

Section 2 introduces both inductive and coinductive infinitary de Bruijn-style term systems and recalls basic operations like lifting and substitution. A short excursion addresses named analogues to the de Bruijn-systems and reviews the notions of α-equality and bound variable renaming, without going into detail. Finally, the abstract framework of notations for reduction relations and basic closure properties of the general notions like parallel substitutivity are discussed.

Section 3 defines $\Pi\mathsf{IE}$-systems with the core components of β-reduction and permutation, provides some examples and establishes a first result on the compatibility of β-reduction with parallel substitutive reduction relations. For wellfounded $\Pi\mathsf{IE}$-systems, inductive characterizations of terms and normal forms are displayed.

Section 4 is dedicated to the confluence problem for β-reduction. First, the Takahashi proof for the λ-calculus Λ is repeated in order to demonstrate how the notation systems for developments behave. Although the basic definitions carry over to full $\Pi\mathsf{IE}$-systems, this does not suffice to infer confluence for them. A detailed study of coinductive Λ exhibits more counterexamples before the concept of bounding finds its first application to β-developments. Finally, this method is extended to bounded $\Pi\mathsf{IE}^{\mathrm{co}}$-systems, which have to be defined and analyzed.

Section 5 proves confluence for permutations in $\Pi\mathsf{IE}^{\mathrm{co}}$-systems. To this end, the [JM00] proof for $\Lambda\mathsf{J}$ is recast in the framework of reduction notations and extended to $\Pi\mathsf{IE}$-systems as far as possible. This is not very far, but remedy can be found in a study of the coinductive analogue of $\Lambda\mathsf{J}$. Ramified permutative developments for $\Lambda\mathsf{J}^{\mathrm{co}}$ are introduced and bounded, until they disclose confluence for the example calculus. Then the general approach is explained, where an abstract apparatus for bounding and ramification of developments along trees becomes necessary.

Section 6 first treats commutation of β-reduction and permutation in $\Lambda\mathsf{J}$. Then it analyzes two different conciliation properties for the reduction relations in general $\Pi\mathsf{IE}^{\mathrm{co}}$-systems and proves commutation relative to the weaker one, again with the help of a bounding argument.

Section 7 focuses on standardization results for $\Pi\mathsf{IE}$-systems. The basic