Strict Ideal Completions of the Lambda Calculus

Patrick Bahr

IT University of Copenhagen, Denmark paba@itu.dk https://orcid.org/0000-0003-1600-8261

Abstract

The infinitary lambda calculi pioneered by Kennaway et al. extend the basic lambda calculus by metric completion to infinite terms and reductions. Depending on the chosen metric, the resulting infinitary calculi exhibit different notions of *strictness*. To obtain infinitary normalisation and infinitary confluence properties for these calculi, Kennaway et al. extend β -reduction with infinitely many ' \perp -rules', which contract *meaningless terms* directly to \perp . Three of the resulting Böhm reduction calculi have unique infinitary normal forms corresponding to Böhm-like trees.

In this paper we develop a corresponding theory of infinitary lambda calculi based on ideal completion instead of metric completion. We show that each of our calculi conservatively extends the corresponding metric-based calculus. Three of our calculi are infinitarily normalising and confluent; their unique infinitary normal forms are exactly the Böhm-like trees of the corresponding metric-based calculi. Our calculi dispense with the infinitely many \perp -rules of the metric-based calculi. The fully non-strict calculus (called 111) consists of only β -reduction, while the other two calculi (called 001 and 101) require two additional rules that precisely state their strictness properties: $\lambda x \perp \rightarrow \perp$ (for 001) and $\perp M \rightarrow \perp$ (for 001 and 101).

2012 ACM Subject Classification Theory of computation \rightarrow Rewrite systems

Keywords and phrases lambda calculus, infinitary rewriting, Böhm trees, meaningless terms, confluence

Digital Object Identifier 10.4230/LIPIcs.FSCD.2018.8

Related Version For space considerations we abridged and in some cases omitted proofs. The corresponding full proofs can be found in the extended version of this paper [5], https://arxiv. org/abs/1805.06736.

1 Introduction

In their seminal work on infinitary lambda calculus, Kennaway et al. [9] study different infinitary variants of the lambda calculus, which are obtained by extending the ordinary lambda calculus by means of metric completion. Different variants of the calculus are obtained by choosing a different metric. The 'standard' metric on terms measures the distance between two terms depending on how deep one has to go into the term structure to distinguish two terms. For example the term xy is closer to the term xz than to the term x, because in the former case both terms are applications whereas in the latter case one term is an application and the other is a variable.

The different metric spaces arise by changing the way in which we measure depth. Kennaway et al. [9] indicate this using a binary triple abc with $a, b, c \in \{0, 1\}$, where a = 0indicates that we do not count lambda abstractions when calculating the depth, and b = 0or c = 0 indicates that we do not count the left or the right side of applications, respectively. More intuitively these three parameters can be interpreted as indicating strictness. For example, a = 0 indicates that lambda abstraction is strict, i.e. if M diverges, then so does $\lambda x.M.$

© Patrick Bahr \odot

licensed under Creative Commons License CC-BY

3rd International Conference on Formal Structures for Computation and Deduction (FSCD 2018). Editor: Hélène Kirchner; Article No. 8; pp. 8:1-8:16



Leibniz International Proceedings in Informatics LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

> The official version will be available from July 9, 2018 at: http://www.dagstuhl.de/dagpub/978-3-95977-077-4

8:2 Strict Ideal Completions of the Lambda Calculus

Since the set of infinite terms is constructed from the set of finite terms by means of metric completion, each calculus has a different universe of terms, as well as a different mode of convergence, which is based on the topology induced by the metric. For instance, from the lambda term $N = (\lambda x.x x y)(\lambda x.x x y)$, we can derive the infinite reduction $N \to N y \to N y y \to \dots$ In the fully non-strict calculus, where abc = 111, this reduction converges to the infinite term $M = \dots y y y$ (i.e. M satisfies M = M y). By contrast, in the calculus 101, which is strict on the left-hand side of every application, this reduction does not converge. In fact, M is not even a valid term in the 101 calculus.

In order to deal with divergence as exemplified for the 101 calculus above, Kennaway et al. [9] extend standard β -reduction to Böhm reduction by adding rules of the form $M \to \bot$, for each term M that causes divergence such as the term N in the 101 calculus. The resulting 001, 101, and 111 calculi based on Böhm reduction have unique normal forms, which correspond to the well-known Böhm Trees [14, 6], Levy-Longo Trees [12, 13] and Berarducci Trees [7], respectively.

In this paper, we introduce infinitary lambda calculi that are based on ideal completion instead of metric completion with the goal of directly dealing with diverging terms without the need for additional reduction rules that contract diverging terms immediately to \perp . To this end, we devise for each metric of the calculi of Kennaway et al. [9] a corresponding partial order with the following property: Ideal completion of the set of finite lambda terms yields the same set of infinite lambda terms as the corresponding metric completion (Section 3). We also find a strong correspondence between the modes of convergence induced by these structures: Each ideal completion yields a complete semilattice structure, which means that the *limit inferior* is always defined. We show that this limit inferior is a conservative extension of the limit in the corresponding metric completion in the sense that both modes of convergence coincide on total lambda terms, i.e. terms without \perp (Section 3).

Based on these partial order structures we define infinitary lambda calculi by a straightforward instantiation of transfinite abstract reduction systems [2]. We find that the ideal completion calculi form a conservative extension of the metric completion calculi of Kennaway et al. [9] (Section 4). Moreover, in analogy to Blom [16] and Bahr [3], we find that the differences between the ideal completion approach and the metric completion approach are compensated for by adding \perp -rules to the metric calculi in the style of Kennaway et al. [10] (Section 5). Finally, we also show infinitary normalisation for our ideal completion calculi and infinitary confluence for the 001, 101, and 111 calculi (Section 5). However, in order to obtain infinitary confluence for 001 and 101, we need to extend β -reduction with two additional rules that precisely capture the strictness properties of these calculi: $\lambda x. \perp \rightarrow \perp$ (for 001) and $\perp M \rightarrow \perp$ (for 001 and 101). In Section 6, we give a brief overview of related work.

2 The Metric Completion

In this section, we introduce infinite lambda terms as the result of metric completion of the set of finite lambda terms. Before we get started, we introduce some basic notions about transfinite sequences and lambda terms. We presume basic familiarity with metric spaces and ordinal numbers.

A sequence over a set A of length α is a mapping from an ordinal α into A and is written as $(a_{\iota})_{\iota < \alpha}$, which indicates the mapping $\iota \mapsto a_{\iota}$; the notation $|(a_{\iota})_{\iota < \alpha}|$ denotes the length α of $(a_{\iota})_{\iota < \alpha}$. If α is a limit ordinal, then $(a_{\iota})_{\iota < \alpha}$ is called *open*; otherwise it is called *closed*. If $(a_{\iota})_{\iota < \alpha}$ is finite, it is also written as $\langle a_0, \ldots, a_{\alpha-1} \rangle$; in particular, $\langle \rangle$ denotes the empty sequence. We write $S \cdot T$ for the *concatenation* of two sequences S and T; S is called a (*proper*) prefix of T, denoted $S \leq T$ (resp. S < T) if there is a (non-empty) sequence S' such that $S \cdot S' = T$. The unique prefix of a sequence S of length $\beta \leq |S|$ is denoted by $S|_{\beta}$.

We consider lambda terms with an additional symbol \perp ; the resulting set of *lambda terms* Λ_{\perp} is inductively defined by the following grammar:

$$M, N ::= \perp \mid x \mid \lambda x.M \mid MN$$

where x is drawn from a countably infinite set \mathcal{V} of variable symbols. The set of *total lambda* terms Λ is the subset of lambda terms in Λ_{\perp} that do not contain \perp . Occurrences of a variable x in a subterm $\lambda x.M$ are called *bound*; other occurrences are called *free*. We use the notation $M[x \to y]$ to replace all free occurrences of the variable x in M with the variable y. We use finite sequences over $\{0, 1, 2\}$, called *positions*, to point to subterms of a lambda term; we write \mathcal{P} for the set of all positions. For each $M \in \Lambda_{\perp}$, $\mathcal{P}(M)$ denotes the set of positions of M (excluding ' \perp 's) recursively defined as follows: $\mathcal{P}(\perp) = \emptyset$, $\mathcal{P}(x) = \{\langle \rangle\}$, $\mathcal{P}(M_1 M_2) =$ $\{\langle \rangle\} \cup \{\langle i \rangle \cdot p \mid i \in \{1, 2\}, p \in \mathcal{P}(M_i)\}$, and $\mathcal{P}(\lambda x.M) = \{\langle \rangle\} \cup \{\langle 0 \rangle \cdot p \mid p \in \mathcal{P}(M)\}$.

A conflict [9] between two lambda terms M, N is a position $p \in \mathcal{P}(M) \cup \mathcal{P}(N)$ such that: (a) if $p = \langle \rangle$, then M and N are not identical variables, not both \bot , not both applications, and not both abstractions; (b) if $p = \langle i \rangle \cdot q$ and $i \in \{1, 2\}$, then $M = M_1 M_2$, $N = N_1 N_2$, and q is a conflict of M_i and N_i ; (c) if $p = \langle 0 \rangle \cdot q$, then $M = \lambda x.M'$, $N = \lambda y.N'$, and q is a conflict of $M'[x \to z]$ and $N'[y \to z]$, where z is a fresh variable occurring neither in M nor N. The terms M and N are said to be α -equivalent if they have no conflicts. By convention we identify α -equivalent terms (i.e. Λ_{\perp} and Λ are assumed to be quotients by α -equivalence).

▶ **Definition 2.1.** Given a triple $\overline{a} = a_0 a_1 a_2 \in \{0, 1\}^3$, called *strictness signature*, a position is called \overline{a} -strict if it is of the form $q \cdot \langle i \rangle$ with $a_i = 0$; otherwise it is called \overline{a} -non-strict. If \overline{a} is clear from the context, we only say strict resp. non-strict.

That is, a strictness signature indicates strictness by 0 and non-strictness by 1. For example, if $\overline{a} = 011$, lambda abstraction is strict, and application is non-strict both from the left and the right. We shall see what this means shortly: Following Kennaway et al. [9], we derive, from a strictness signature \overline{a} , a depth measure $\left|\cdot\right|^{\overline{a}}$, which counts the number of non-strict, non-empty prefixes of a position. From this depth measure we then derive a corresponding metric $\mathbf{d}^{\overline{a}}$ on lambda terms.

▶ **Definition 2.2.** Given a strictness signature \overline{a} , the \overline{a} -depth of a position p, denoted $|p|^{\overline{a}}$, is recursively defined as $|\langle\rangle|^{\overline{a}} = 0$ and $|q \cdot \langle i \rangle|^{\overline{a}} = |q|^{\overline{a}} + a_i$. The \overline{a} -distance $\mathbf{d}^{\overline{a}}(M, N)$ between two terms $M, N \in \Lambda_{\perp}$ is 0 if M and N are α -equivalent and otherwise 2^{-d} , where d is the least number satisfying $d = |p|^{\overline{a}}$ for some conflict p of M and N.

Kennaway et al. [9] showed that the pair $(\Lambda_{\perp}, \mathbf{d}^{\overline{a}})$ forms an ultrametric space for any \overline{a} . Intuitively, the consequence of the definition of these metric spaces is that sequences of terms, such as the sequence $N, N y, N y y, \ldots$, only converge if conflicts between consecutive terms are guarded by an increasing number of non-strict positions. In the example, conflicts between consecutive terms are guarded by an increasing stack of applications to y. If $a_1 = 1$, these applications correspond to non-strict positions, and thus the sequence converges. However, if $a_1 = 0$, the sequence does not converge.

We turn now to the metric completion. To facilitate later definitions and to illustrate the resulting structures, we use a partial function representation in the form of lambda trees taken from Blom [16], which will serve as mediator between metric completion and ideal

8:4 Strict Ideal Completions of the Lambda Calculus

completion.¹ A lambda tree is a (possibly infinite) labelled tree where a label λ indicates abstraction and @ indicates application; labels in \mathcal{V} indicate free variables and a label $p \in \mathcal{P}$ indicates a variable that is bound by an abstraction at position p. There is no label corresponding to \bot , which instead is represented as a 'hole' in the tree. We write $\mathcal{D}(f)$ to denote the domain of a partial function f, and $f(p) \simeq g(q)$ to indicate that the partial functions f and q are either both undefined or have the same value at p and q, respectively.

▶ Definition 2.3. A lambda tree is a partial function $t: \mathcal{P} \rightharpoonup \mathcal{L}$ with $\mathcal{L} = \{\lambda, @\} \uplus \mathcal{P} \uplus \mathcal{V}$ so that

- $\begin{array}{ll} p \cdot \langle 0 \rangle \in \mathcal{D}(t) & \implies \quad t(p) = \lambda, \\ p \cdot \langle 1 \rangle \in \mathcal{D}(t) \text{ or } p \cdot \langle 2 \rangle \in \mathcal{D}(t) & \implies \quad t(p) = @, \text{ and} \\ t(p) = q, \text{ where } q \in \mathcal{P} & \implies \quad q \leq p \text{ and } t(q) = \lambda. \end{array}$ (a) (b)
- (c)

As one would expect, the domain $\mathcal{D}(t)$ of a lambda tree t is prefix closed.

The set of all lambda trees is denoted $\mathcal{T}_{\perp}^{\infty}$. The set of \perp -positions in t, denoted $\mathcal{D}_{\perp}(t)$, is the smallest set satisfying (a) $\langle \rangle \notin \mathcal{D}(t)$ implies $\langle \rangle \in \mathcal{D}_{\perp}(t)$; (b) $t(p) = \lambda, p \cdot \langle 0 \rangle \notin \mathcal{D}(t)$ implies $p \cdot \langle 0 \rangle \in \mathcal{D}_{\perp}(t)$; and (c) $t(p) = @, p \cdot \langle i \rangle \notin \mathcal{D}(t), i \in \{1, 2\}$ implies $p \cdot \langle i \rangle \in \mathcal{D}_{\perp}(t)$. A lambda tree t is called total if $\mathcal{D}_{\perp}(t)$ is empty. The set of all total lambda trees is denoted \mathcal{T}^{∞} . A lambda tree t is called *finite* if $\mathcal{D}(t)$ is a finite set. The set of all finite (total) lambda trees is denoted \mathcal{T}_{\perp} (respectively \mathcal{T}). A renaming of a lambda tree t is a lambda tree s such that there is a bijection $f: \mathcal{V} \to \mathcal{V}$ with the following properties: s(p) = t(p) if $t(p) \in \mathcal{L} \setminus \mathcal{V}$, s(p) = f(t(p)) if $t(p) \in \mathcal{V}$, and otherwise s(p) is undefined.

In order to avoid confusion, we use upper case letters M, N for lambda terms and lower case letters s, t, u for lambda trees. Below, we give a bijection from lambda terms to finite lambda trees that should help illustrate the idea behind lambda trees. At the heart of this bijection are the following constructions based on Blom [16]:

▶ Definition 2.4. Given lambda trees $t, t_1, t_2 \in \mathcal{T}_{\perp}^{\infty}$ and a variable $x \in \mathcal{V}$, let \perp , x, λ x.t and $t_1 t_2$ be partial functions of type $\mathcal{P} \rightarrow \mathcal{L}$ defined by their graph as follows:

$$\begin{split} \bot &= \emptyset \qquad \mathbf{x} = \{(\langle \rangle, x)\} \\ \lambda \mathbf{x}.t &= \{(\langle \rangle, \lambda)\} \cup \{(\langle 0 \rangle \cdot p, l) \mid l \in \{\lambda, @\} \uplus \mathcal{V} \setminus \{x\}, (p, l) \in t\} \\ &\cup \{(\langle 0 \rangle \cdot p, \langle 0 \rangle \cdot q) \mid q \in \mathcal{P}, (p, q) \in t\} \cup \{(\langle 0 \rangle \cdot p, \langle \rangle) \mid (p, x) \in t\} \\ t_1 t_2 &= \{(\langle \rangle, @)\} \cup \{(\langle i \rangle \cdot p, l) \mid i \in \{1, 2\}, l \in \{\lambda, @\} \uplus \mathcal{V}, (p, l) \in t_i\} \\ &\cup \{(\langle i \rangle \cdot p, \langle i \rangle \cdot q) \mid i \in \{1, 2\}, q \in \mathcal{P}, (p, q) \in t_i\} \end{split}$$

One can easily check that each of the above four constructions yields a lambda tree, where \perp is the empty lambda tree, x the lambda tree consisting of a single free variable x, λ x.t is a lambda abstraction over x with body t, and $t_1 t_2$ is an application of t_1 to t_2 . The following translation of lambda terms to finite lambda trees illustrates the use of these constructions:

▶ Definition 2.5. Let $\llbracket \cdot \rrbracket : \Lambda_{\perp} \to \mathcal{T}_{\perp}$ be defined recursively as follows:

 $\llbracket \bot \rrbracket = \bot \qquad \llbracket \lambda x.M \rrbracket = \lambda x.\llbracket M \rrbracket$ $\llbracket x \rrbracket = \mathsf{x} \qquad \llbracket M \, N \rrbracket = \llbracket M \rrbracket \ \llbracket N \rrbracket$

One can easily check that $\llbracket \cdot \rrbracket : \Lambda_{\perp} \to \mathcal{T}_{\perp}$ is indeed a bijection, which, if restricted to Λ , is a bijection from Λ to \mathcal{T} . Moreover, one can show that each $t \in \mathcal{T}^{\infty}_{\perp}$ with some $\langle i \rangle \cdot p \in \mathcal{D}(t)$

In the companion report [5] we give a direct proof of the correspondence between metric and ideal completion based on the meta theory of Majster-Cederbaum and Baier [15].

P. Bahr

is equal to $\lambda x.t'$ if i = 0 and to $t_1 t_2$ if $i \in \{1, 2\}$, for some $t', t_1, t_2 \in \mathcal{T}_{\perp}^{\infty}$. Following this observation, we define, for each $t \in \mathcal{T}_{\perp}^{\infty}$ and $p \in \mathcal{D}(t)$, the subtree of t at p, denoted $t|_p$, by induction on p as follows: $t|_{\langle i \rangle} = t$, $\lambda x.t|_{\langle 0 \rangle \cdot p} = t|_p$, and $t_1 t_2|_{\langle i \rangle \cdot p} = t_i|_p$ for $i \in \{1, 2\}$. One can easily check that $t|_p$ is uniquely defined modulo renaming of free variables.

▶ **Definition 2.6.** An *infinite branch* in a lambda tree $t \in \mathcal{T}_{\perp}^{\infty}$ is an infinite sequence S such that each proper prefix of S is in $\mathcal{D}(t)$. We call a proper prefix of S a *position along* S.

Note that by instantiating König's Lemma to lambda trees, we know that a lambda tree is infinite iff it has an infinite branch.

The idea of the metric $\mathbf{d}^{\overline{a}}$ on lambda terms is to disallow (in the ensuing metric completion) infinite branches that have only finitely many non-strict positions along them. The following definition makes this restriction explicit on lambda trees:

▶ **Definition 2.7.** An infinite branch *S* of a lambda tree *t* is called \overline{a} -bounded if the \overline{a} -depth of all positions along *S* is bounded by some $n < \omega$, i.e. $|p|^{\overline{a}} < n$ for all p < S. The lambda tree *t* is called \overline{a} -unguarded if it has an \overline{a} -bounded infinite branch *S*. Otherwise, *t* is called \overline{a} -guarded. The set of all \overline{a} -guarded (total) lambda trees is denoted $\mathcal{T}_{\perp}^{\overline{a}}$ (respectively $\mathcal{T}^{\overline{a}}$). In particular, $\mathcal{T}_{\perp}^{000} = \mathcal{T}_{\perp}$ and $\mathcal{T}_{\perp}^{111} = \mathcal{T}_{\perp}^{\infty}$.

For example, the lambda tree s with s = sy is 101-unguarded while t with $t = \lambda y.ty$ is 101-guarded as each application is guarded by an abstraction (which is non-strict).

For each strictness signature \overline{a} , we give a metric $\mathbf{d}_{\mathcal{T}}^{\overline{a}}$ on lambda trees that corresponds to the metric $\mathbf{d}^{\overline{a}}$ on lambda terms.

▶ **Definition 2.8.** For each two lambda trees $s, t \in \mathcal{T}_{\perp}^{\infty}$, define $\mathbf{d}_{\mathcal{T}}^{\overline{a}}(s,t) = 0$ if s = t and otherwise $\mathbf{d}_{\mathcal{T}}^{\overline{a}}(s,t) = 2^{-d}$, where d is the least $|p|^{\overline{a}}$ with $s(p) \neq t(p)$.

From the characterisation of the metric completion of $(\Lambda_{\perp}, \mathbf{d}^{\overline{a}})$ from Kennaway et al. [9, Lemma 7] we know that the metric space of \overline{a} -guarded lambda trees $(\mathcal{T}_{\perp}^{\overline{a}}, \mathbf{d}_{\mathcal{T}}^{\overline{a}})$ is indeed the metric completion of $(\Lambda_{\perp}, \mathbf{d}^{\overline{a}})$ with the isometric embedding $[\![\cdot]\!] : \Lambda_{\perp} \to \mathcal{T}_{\perp}$ (cf. the companion report [5]). Analogously, $(\mathcal{T}^{\overline{a}}, \mathbf{d}_{\mathcal{T}}^{\overline{a}})$ is the metric completion of $(\Lambda, \mathbf{d}^{\overline{a}})$.

3 The Ideal Completion

In this section, we present an alternative to the metric completion from Section 2 that is based on a family of partial orders on lambda terms indexed by strictness signatures. In the following we assume basic familiarity with order theory.

▶ **Definition 3.1.** Given a strictness signature \overline{a} , the partial order $\leq_{\perp}^{\overline{a}}$ is the least transitive, reflexive order on Λ_{\perp} satisfying the following for all $M, M', N, N' \in \Lambda_{\perp}$ and $x \in \mathcal{V}$:

For the case that $\overline{a} = 111$, we obtain the partial order \leq_{\perp} that is typically used for ideal completions. This order is fully monotone, i.e. $M \leq_{\perp} M'$ implies $\lambda x.M \leq_{\perp} \lambda x.M'$, $MN \leq_{\perp} M'N$ and $NM \leq_{\perp} NM'$. By contrast, $\leq_{\perp}^{\overline{a}}$ restricts monotonicity of abstraction in case $a_0 = 0$ and of application in case $a_1 = 0$ or $a_2 = 0$. Intuitively, we have $M \leq_{\perp}^{\overline{a}} N$ iff N can be obtained from M by replacing occurrences of \perp in M at non-strict positions with

8:6 Strict Ideal Completions of the Lambda Calculus

arbitrary terms. For example, if $\overline{a} = 001$, then neither $\lambda x \perp \leq_{\perp}^{\overline{a}} \lambda x.x x$ nor $\lambda x \perp x \leq_{\perp}^{\overline{a}} \lambda x.x x$; but we do have that $\lambda x.x \perp \leq_{\perp}^{\overline{a}} \lambda x.x x$.

With this intuition in mind, we translate $\leq_{\perp}^{\overline{a}}$ to a corresponding order $\leq_{\perp}^{\overline{a}}$ on lambda trees as follows:

▶ Definition 3.2. Given lambda trees $s, t \in \mathcal{T}_{\perp}^{\infty}$, we have $s \leq_{\perp}^{\overline{a}} t$ if

 $\begin{array}{ll} \text{(a)} & \mathcal{D}(s) \subseteq \mathcal{D}(t), \\ \text{(b)} & s(p) = t(p) & \text{for all } p \in \mathcal{D}(s), \text{ and} \\ \text{(c)} & p \in \mathcal{D}(s) \implies p \cdot \langle i \rangle \in \mathcal{D}(s) & \text{for all } \overline{a}\text{-strict positions } p \cdot \langle i \rangle \in \mathcal{D}(t). \end{array}$

Conditions (a) and (b) alone would give us the corresponding order for the standard partial order \leq_{\perp} . Condition (c) ensures that the partial order $\trianglelefteq_{\perp}^{\overline{a}}$ may not fill a hole in a strict position in the left-hand side tree.

One can check that $(\mathcal{T}_{\perp}^{\infty}, \trianglelefteq_{\perp}^{\overline{a}})$ forms a partially ordered set. Moreover, we have the following correspondence between the two families of orders $\leq_{\perp}^{\overline{a}}$ and $\trianglelefteq_{\perp}^{\overline{a}}$:

▶ Proposition 3.3. $\llbracket \cdot \rrbracket : (\Lambda_{\perp}, \leq_{\perp}^{\overline{a}}) \to (\mathcal{T}_{\perp}, \leq_{\perp}^{\overline{a}})$ is an order isomorphism.

For the remainder of this section, we turn our focus to the partial orders $\trianglelefteq_{\perp}^{\overline{a}}$ on lambda trees. In particular, we show that $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$ forms a *complete semilattice* and that it is (order isomorphic to) the ideal completion of $(\Lambda_{\perp}, \leq_{\perp}^{\overline{a}})$. A complete semilattice is a partially ordered set (A, \leq) that is a *complete partial order* (*cpo*) and that has a *greatest lower bound* (*glb*) $\square B$ for every *non-empty* set $B \subseteq A^2$. A partially ordered set (A, \leq) is a cpo if it has a least element, and each directed set D in (A, \leq) has a *least upper bound* (*lub*) $\bigsqcup D$; a set $D \subseteq A$ is called directed if for each two $a, b \in D$ there is some $c \in D$ with $a, b \leq c$.

In particular, for any sequence $(a_{\iota})_{\iota < \alpha}$ in a complete semilattice, its *limit inferior*, defined by $\liminf_{\iota \to \alpha} a_{\iota} = \bigsqcup_{\beta < \alpha} \left(\bigcap_{\beta \leq \iota < \alpha} a_{\iota} \right)$, exists. While the metric completion lambda calculi are based on the limit of the underlying metric space, our ideal completion lambda calculi are based on the limit inferior.

To show that $(\mathcal{T}^{\overline{a}}_{\perp}, \leq^{\overline{a}}_{\perp})$ forms a complete semilattice structure, we construct the appropriate lubs and glbs:

▶ **Theorem 3.4 (cpo** $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$). The partially ordered set $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$ forms a complete partial order. In particular, the lub t of a directed set D satisfies the following:

 $\mathcal{D}(t) = \bigcup_{s \in D} \mathcal{D}(s) \qquad \qquad s(p) = t(p) \quad \text{for all } s \in D, p \in \mathcal{D}(s)$

Proof sketch. The lambda tree \perp is the least element in $(\mathcal{T}_{\perp}^{\overline{a}}, \leq \underline{a}_{\perp})$. Construct the lub t of D as follows: t(p) = s(p) iff there is some $s \in D$ with $p \in \mathcal{D}(s)$. One can check that t indeed is a well-defined lambda tree that is \overline{a} -guarded and is the least upper bound of D.

▶ **Proposition 3.5** (glbs of $\trianglelefteq_{\perp}^{\overline{a}}$). Every non-empty subset T of $\mathcal{T}_{\perp}^{\overline{a}}$ has a glb $\prod T$ in $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$ such that $\mathcal{D}(\prod T)$ is the largest set P satisfying the following properties:

- (1) If $p \in P$, then there is some $l \in \mathcal{L}$ such that s(p) = l for all $s \in T$.
- (2) If $p \cdot \langle i \rangle \in P$, then $p \in P$.

(3) If $p \in P$, $a_i = 0$, and $p \cdot \langle i \rangle \in \mathcal{D}(s)$ for some $s \in T$, then $p \cdot \langle i \rangle \in P$.

² Equivalently, complete semilattices are bounded complete cpos. Hence, complete semilattices are a generalisation of *Scott domains* (which in addition have to be *algebraic*).

P. Bahr

Proof sketch. Let $P \subseteq \mathcal{P}$ be the largest set satisfying (1) to (3). As these properties are closed under union, P is well-defined. We define the partial function $t: \mathcal{P} \to \mathcal{L}$ as the restriction of an arbitrary lambda tree in T to P. Using (1) and (2), one can show that t is indeed a well-defined \overline{a} -guarded lambda tree. One can then check that t is the glb of T.

For instance $\prod \{\lambda x. x y, \lambda x. y x\}$ is $\lambda x. \perp \perp$ for 011, $\lambda x. \perp$ for 110, and \perp for 001.

▶ Theorem 3.6. $(\mathcal{T}_{\perp}^{\overline{a}}, \leq \overline{a}_{\perp})$ is a complete semilattice for any \overline{a} .

Proof. Follows from Theorem 3.4 and Proposition 3.5.

We conclude this section by establishing the partially ordered set $(\mathcal{T}_{\perp}^{\overline{a}}, \leq \underline{a}_{\perp}^{\overline{a}})$ as (order isomorphic to) the ideal completion of $(\Lambda_{\perp}, \leq \underline{a}_{\perp}^{\overline{a}})$. Recall that, given a partially order set (A, \leq) , its ideal completion is an extension of the original partially ordered set to a cpo. A set $B \subseteq A$ is called an *ideal* in (A, \leq) if it is directed and *downward-closed*, where the latter means that for all $a \in A, b \in B$ with $a \leq b$, we have that $a \in B$. The *ideal completion* of (A, \leq) , is the partially ordered set (\mathcal{I}, \subseteq) , where \mathcal{I} is the set of all ideals in (A, \leq) and \subseteq is standard set inclusion.

▶ Theorem 3.7. The ideal completion of $(\Lambda_{\perp}, \leq_{\perp}^{\overline{a}})$ is order isomorphic to $(\mathcal{T}_{\perp}^{\overline{a}}, \leq_{\perp}^{\overline{a}})$.

Proof sketch. By Proposition 3.3, it suffices to show that the ideal completion (\mathcal{I}, \subseteq) of $(\mathcal{T}_{\perp}, \trianglelefteq_{\perp}^{\overline{a}})$ is order isomorphic to $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$. To this end, we define two functions $\phi: \mathcal{T}_{\perp}^{\overline{a}} \to \mathcal{I}$ and $\psi: \mathcal{I} \to \mathcal{T}_{\perp}^{\overline{a}}$ as follows: $\phi(t) = \{s \in \mathcal{T}_{\perp} \mid s \trianglelefteq_{\perp}^{\overline{a}} t\}; \psi(\mathcal{I}) = \bigsqcup T$. Well-definedness of ϕ and ψ follows from König's Lemma and Theorem 3.4, respectively. Both ϕ and ψ are obviously monotone and one can check that ϕ and ψ are inverses of each other. Hence, (\mathcal{I}, \subseteq) is order isomorphic to $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$

Now that we have established the connection between $\mathcal{T}_{\perp}^{\overline{a}}$ and the metric completion resp. the ideal completion of Λ_{\perp} , we turn our focus to $\mathcal{T}_{\perp}^{\overline{a}}$ for the rest of this paper.

The characterisation of lubs and glbs for the complete semilattice $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$ allows us to relate the corresponding notion of limit inferior with the limit in the complete metric space $(\mathcal{T}_{\perp}^{\overline{a}}, \mathbf{d}_{\mathcal{T}}^{\overline{a}})$ as summarised in the following theorem:

▶ Theorem 3.8. Let $(t_{\iota})_{\iota < \alpha}$ be a sequence in $\mathcal{T}_{\downarrow}^{\overline{a}}$.

- (i) If $\lim_{\iota \to \alpha} t_{\iota} = t$ in $(\mathcal{T}_{\perp}^{\overline{a}}, \mathbf{d}_{\mathcal{T}}^{\overline{a}})$, then $\liminf_{\iota \to \alpha} t_{\iota} = t$ in $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{\overline{a}})$.
- (ii) If $\liminf_{\iota \to \alpha} t_{\iota} = t$ in $(\mathcal{T}_{\perp}^{\overline{a}}, \leq \underline{a}_{\perp}^{\overline{a}})$ and t is total, then $\lim_{\iota \to \alpha} t_{\iota} = t$ in $(\mathcal{T}_{\perp}^{\overline{a}}, \mathbf{d}_{\mathcal{T}}^{\overline{a}})$.

The key to establish the correspondence above is the following characterisation of the limit t of a converging sequence $(t_{\iota})_{\iota < \alpha}$ in $(\mathcal{T}^{\overline{a}}_{\perp}, \mathbf{d}^{\overline{a}}_{\mathcal{T}})$:

$$\mathcal{D}(t) = \bigcup_{\beta \leq \alpha} \bigcap_{\beta \leq \iota \leq \alpha} \mathcal{D}(t_{\iota}), \text{ and } t(p) = l \iff \exists \beta < \alpha \forall \beta \leq \iota < \alpha \colon t_{\iota}(p) = l$$

The proof of the correspondence result makes use of a notion of truncation similar Arnold and Nivat's [1] but generalised to be compatible with the $\leq \frac{\overline{a}}{1}$ -orderings.

From the above findings we can conclude that the limit inferior in $(\mathcal{T}_{\perp}^{\overline{a}}, \leq \underline{a}_{\perp}^{\overline{a}})$ restricted to total lambda trees coincides with the limit in $(\mathcal{T}^{\overline{a}}, \mathbf{d}_{\mathcal{T}}^{\overline{a}})$. In other words, the limit inferior is a conservative extension of the limit. In the next section, we transfer this result to (strong) convergence of reductions.

8:8 Strict Ideal Completions of the Lambda Calculus

4 Transfinite Reductions

In this section, we study finite and transfinite reductions on lambda trees. To this end, we assume for the remainder of this paper a fixed strictness signature \bar{a} such that all subsequent definitions and theorems work on the same universe of lambda trees $\mathcal{T}_{\perp}^{\bar{a}}$ and its associated structures $\mathbf{d}_{\mathcal{T}}^{\bar{a}}$ and $\leq_{\perp}^{\bar{a}}$ (unless stated otherwise). Moreover, we need a suitably general notion of reduction steps beyond the familiar β - and η -rules in order to accommodate Böhm reductions in Section 5.

▶ **Definition 4.1.** A rewrite system R is a binary relation on $\mathcal{T}_{\perp}^{\overline{a}}$ such that $(s,t) \in R$ implies that $s \neq \perp$. Given $s, t \in \mathcal{T}_{\perp}^{\overline{a}}$ and $p \in \mathcal{P}$, an *R*-reduction step from s to t at p, denoted $s \rightarrow_{R,p} t$, is inductively defined as follows: if $(s,t) \in R$, then $s \rightarrow_{R,\langle\rangle} t$; if $t \rightarrow_{R,p} t'$, then $\lambda x.t \rightarrow_{R,\langle0\rangle \cdot p} \lambda x.t', t s \rightarrow_{R,\langle1\rangle \cdot p} t's$, and $st \rightarrow_{R,\langle2\rangle \cdot p} st'$ for all $s \in \mathcal{T}_{\perp}^{\overline{a}}$. If R or p are irrelevant or clear from the context, we omit them in the notation $\rightarrow_{R,p}$. If $(t,t') \in R$, then t is called an *R*-redex. If $s \rightarrow_{R,p} t$, then s is said to have an *R*-redex occurrence at p. A lambda tree t is called an *R*-normal form if no *R*-reduction step starts from t. The prefix "*R*-" is dropped if R is irrelevant or clear from the context.

Example 4.2. The familiar β - and η -rules form rewrite systems as follows:

$$\beta = \left\{ \left((\lambda \mathbf{x}.t) \, s, t \, [x/s] \right) \, \middle| \, s, t \in \mathcal{T}_{\perp}^{\overline{a}} \right\} \qquad \text{in} = \left\{ \left(\lambda \mathbf{x}.t \, \mathbf{x}, t \right) \, \middle| \, t \in \mathcal{T}_{\perp}^{\overline{a}}, x \notin \mathsf{Range}(t) \right\}$$

where substitution t[x/s] is defined as follows: for each $p \in \mathcal{P}$ we have t[x/s](p) = t(p) if $t(p) \in \mathcal{L} \setminus \{x\}; t[x/s](p) = s(p_2)$ if $p = p_1 \cdot p_2, t(p_1) = x, s(p_2) \in \mathcal{L} \setminus \mathcal{P}; t[x/s](p) = p_1 \cdot s(p_2)$ if $p = p_1 \cdot p_2, t(p_1) = x, s(p_2) \in \mathcal{P};$ and t[x/s](p) is undefined otherwise.

The resulting β -reduction step relation \rightarrow_{β} on lambda trees is isomorphic (via the isomorphism of Theorem 3.7) to the lifting of the ordinary finitary β -reduction step relation on lambda terms to the ideal completion via the lifting operator [\cdot) of Blom [8]. An analogous correspondence can be shown for η as well.

▶ **Definition 4.3.** A sequence $S = (t_{\iota} \rightarrow_{R,\rho_{\iota}} t_{\iota+1})_{\iota < \alpha}$ of *R*-reduction steps is called an *R*-reduction; *S* is called *total* if each t_{ι} is total. If *S* is finite, we also write $S: t_0 \rightarrow_R^* t_{\alpha}$.

The above notion of reductions is too general as it does not relate lambda trees t_{β} at a limit ordinal index β to the lambda trees $(t_{\iota})_{\iota < \beta}$ that precede it. This shortcoming is addressed with a suitable notion of convergence and continuity. In the literature on infinitary rewriting one finds two different variants of convergence/continuity: a *weak* variant, which defines convergence/continuity only according to the underlying structure (metric limit or limit inferior), and a *strong* variant, which also takes the position of contracted redexes into consideration. While both the metric and the partial order lend themselves to either variant, we only consider the strong variant here and refer the reader to the companion report [5] for the weak variant.

We use the name m-convergence and p-convergence to distinguish between the metricand the partial order-based notion of convergence, respectively. Our notion of (strong) mconvergence is the same notion of convergence that Kennaway et al. [9] used for their infinitary lambda calculi. For our notion of (strong) p-convergence we instantiate the abstract notion of strong p-convergence from our previous work [2]. The key ingredient of p-convergence is the notion of *reduction context*, which assigns to each reduction step $s \to t$ a lambda tree cwith $c \leq \underline{a}_{\perp}^{\overline{a}} s, t$. Intuitively, this reduction context c comprises the (maximal) fragment of sthat cannot be changed by the reduction step, regardless of the reduction rule. For instance, the reduction context of $\lambda x.(\lambda y.y) \times \to \lambda x.x$ is $\lambda x.\perp$ if $a_0 = 1$, and \perp otherwise. The notion of p-convergence is defined using the limit inferior of the sequence of reduction contexts (instead of the original lambda trees themselves). The canonical approach to derive such a reduction context for any complete semilattice is to take the greatest lower bound of the involved lambda trees s and t that does not contain any position of the redex:

▶ **Definition 4.4.** The reduction context of a reduction step $s \to_p t$ is the greatest lambda tree c in $(\mathcal{T}^{\overline{a}}_{\perp}, \trianglelefteq^{\overline{a}}_{\perp})$ with $c \trianglelefteq^{\overline{a}}_{\perp} s, t$ and $p \notin \mathcal{D}(c)$; we write $s \to_c t$ to indicate the reduction context c.

In order to simplify reasoning and provide an intuitive understanding of the concept, we give a direct construction of reduction contexts as well:

▶ **Definition 4.5.** Given $t \in \mathcal{T}_{\perp}^{\infty}$ and $p \in \mathcal{D}(t)$, we write $t \setminus p$ for the restriction of t to the domain $\{q \in \mathcal{D}(t) \mid p \leq q\}$, and $p \downarrow^{\overline{a}}$ for the longest non-strict prefix of p.

That is, $t \setminus p$ is obtained from t by replacing the subtree at p with \bot . Moreover, $\downarrow^{\overline{a}}$ can be characterised as follows: $\langle \rangle \downarrow^{\overline{a}} = \langle \rangle$; $(p \cdot \langle i \rangle) \downarrow^{\overline{a}} = p \cdot \langle i \rangle$ if $a_i = 1$; and $(p \cdot \langle i \rangle) \downarrow^{\overline{a}} = p \downarrow^{\overline{a}}$ if $a_i = 0$.

▶ Lemma 4.6. The reduction context of $s \rightarrow_p t$ is equal to $s \rightarrow_p \downarrow^{\overline{a}}$ and $t \setminus p \downarrow^{\overline{a}}$.

Proof sketch. By a straightforward induction on p.

That is, the reduction context of $s \rightarrow_p t$ is obtained from s by removing the most deeply nested subtree that both contains the redex and is in a non-strict position. The ensuing notions of strong convergence of reductions are spelled out as follows:

▶ **Definition 4.7.** An *R*-reduction $S = (t_{\iota} \rightarrow_{p_{\iota},c_{\iota}} t_{\iota+1})_{\iota < \alpha}$ m-converges to t_{α} , denoted $S: t_0 \xrightarrow{\mathsf{m}}_{R} t_{\alpha}$, if $\lim_{\iota \to \gamma} t_{\iota} = t_{\gamma}$ and $(p_{\iota})^{\overline{\mathfrak{a}}}_{\iota < \gamma}$ tends to infinity for all limit ordinals $\gamma \leq \alpha$. S p-converges to t_{α} , denoted $S: t_0 \xrightarrow{\mathsf{p}}_{R} t_{\alpha}$, if $\liminf_{\iota \to \gamma} c_{\iota} = t_{\gamma}$ for all limit ordinals $\gamma \leq \alpha$. S is called m-continuous resp. p-continuous if the corresponding convergence conditions hold for limit ordinals $\gamma < \alpha$ (instead of $\gamma \leq \alpha$).

Intuitively, strong convergence under-approximates convergence in the underlying structure (i.e. weak convergence) by assuming that every contraction changes the root symbol of the redex. Thus, given a reduction step $s \rightarrow_p t$, strong convergence assumes that the shortest position at which s and t differ is p.

The semilattice structure underlying p-convergence ensures that p-continuous reductions always p-converge, whereas m-convergence does not necessarily follow from m-continuity:

► Example 4.8. Given $\Omega = (\lambda x. x x)(\lambda x. x x)$ and $t = (\lambda x. x \Omega) y$, we consider the β -reduction $S: t \to t \to \ldots$ that repeatedly contracts the redex Ω in t. S is trivially m- and p-continuous. However, it is not m-convergent, since contraction takes place at a constant \overline{a} -depth, namely $|\langle 1, 0, 2 \rangle|^{\overline{a}}$. But it p-converges to $t \setminus \langle 1, 0, 2 \rangle \downarrow^{\overline{a}}$, which is also the reduction context of each reduction step in S and is equal to $(\lambda x. x \bot) y$ if $a_2 = 1$, to $(\lambda x. \bot) y$ if $a_2 = 0$ but $a_0 = 1$, to $\bot y$ if $\overline{a} = 010$, and to \bot if $\overline{a} = 000$.

Similarly to the correspondence between the limit and the limit inferior in Theorem 3.8, we find a correspondence between p- and m-convergence.

▶ **Proposition 4.9.** For each reduction $S: s \xrightarrow{\mathsf{m}} t$, we also have that $S: s \xrightarrow{\mathsf{p}} t$.

Proof sketch. Let $S = (t_{\iota} \to_{p_{\iota},c_{\iota}} t_{\iota+1})_{\iota < \alpha}$. If S m-converges, then $(|p_{\iota}|^{\overline{\alpha}})_{\iota < \gamma}$ tends to infinity for all limit ordinals $\gamma < \alpha$, i.e. for each $d < \omega$ we have that $|p_{\iota}|^{\overline{\alpha}} \ge d$ after some $\delta < \gamma$. With the help of Lemma 4.6, one can show that the latter implies that t_{ι} and c_{ι} coincide up to

FSCD 2018

8:10 Strict Ideal Completions of the Lambda Calculus

 \overline{a} -depth d for all $\delta \leq \iota < \gamma$. Consequently, $\lim_{\iota \to \gamma} t_{\iota} = \lim_{\iota \to \gamma} c_{\iota}$, which, by Theorem 3.8 (i), implies $\lim_{\iota \to \gamma} t_{\iota} = \liminf_{\iota \to \gamma} c_{\iota}$. Since this holds for all limit ordinals $\gamma \leq \alpha$, we know that S also p-converges to t.

With the proposition above, we derive the other direction of the correspondence:

▶ **Proposition 4.10.** $S: s \xrightarrow{p} t$ implies $S: s \xrightarrow{m} t$ whenever S and t are total.

Proof sketch. One can show that the totality of S and t implies that the \overline{a} -depth of contracted redexes in each open prefix of S tends to infinity. Using Proposition 5.5 from [2], we can show that the latter implies that S also m-converges. Then according to Proposition 4.9, S must m-converge to the same lambda tree t.

Note that it is not sufficient that the two trees s and t are total. For example, the β -reduction $S: (\lambda x.y) \Omega \xrightarrow{\mathbf{p}} (\lambda x.y) \perp \rightarrow y$ p-converges to y but does not m-converge.

Putting Propositions 4.9 and 4.10 together we obtain that p-convergence is a conservative extension of m-convergence:

▶ Corollary 4.11. $S: s \xrightarrow{\mathsf{m}} t$ iff $S: s \xrightarrow{\mathsf{p}} t$ whenever S and t are total.

5 Beta Reduction

So far we have only studied the properties of p-convergence independent of the rewrite system. In this section, we specifically study β -reduction and show infinitary normalisation for all of our calculi, and infinitary confluence for three of them. However, considering pure β -reduction, infinitary confluence only holds for the 111 calculus. We can construct counterexamples for the other calculi:

► Example 5.1 ([9]). Given $a_2 = 0$ and $t = (\lambda x. y) \Omega$, we find reductions $t \xrightarrow{\mathbb{P}}_{\mathfrak{P}} \bot$ and $t \to_{\mathfrak{G}} y$. Given $a_2 = 1$, $a_1 = 0$, and $t = (\lambda x. x y) \Omega$, we have $t \xrightarrow{\mathbb{P}}_{\mathfrak{P}} (\lambda x. x y) \bot \to_{\mathfrak{G}} \bot y$ and $t \to_{\mathfrak{G}} \Omega y \xrightarrow{\mathbb{P}}_{\mathfrak{P}} \bot$. Similarly, given $a_2 = 1$, $a_0 = 0$, and $t = (\lambda x. \lambda y. x) \Omega$, we have $t \xrightarrow{\mathbb{P}}_{\mathfrak{P}} (\lambda x. \lambda y. x) \bot \to_{\mathfrak{G}} \lambda y. \bot \to_{\mathfrak{G}} \lambda y. \bot$ and $t \to_{\mathfrak{G}} \lambda y. \Omega$.

Infinitary confluence of pure β -reduction fails for all m-convergence calculi of Kennaway et al.[9] – including the 111 calculus. On the other hand, the Böhm reduction calculi of Kennaway et al. [10], which extend pure β -reduction with infinitely many rules of the form $t \to \perp$, do satisfy infinitary confluence for the 001, 101, and 111 calculi.

We would like to obtain similar confluence results for the 001, 101, and 111 p-convergence calculi. However, the gap we have to bridge to achieve infinitary confluence is much narrower in our p-convergence calculi. Intuitively, confluence fails for 001 and 101 because p-convergence only captures partiality that is due to infinite reductions, but not partiality that can propagate via finite reductions: For example, in the 101 calculus we have $\Omega y \stackrel{p_{m_{\beta}}}{\to} \pm$ but $\pm y \stackrel{p_{m_{\beta}}}{\to} \pm$. In order to obtain the desired confluence properties, we have to add the rules $\lambda x. \perp \to \perp$ (for 001) and $\pm t \to \perp$ (for 001 and 101). More generally we define these S-rules formally as follows:

$$\mathbb{S} = \left\{ \left(t_1 \, t_2, \bot\right) \, \middle| \, t_1, t_2 \in \mathcal{T}_{\bot}^{\overline{a}}, t_i = \bot, a_i = 0 \right\} \cup \left\{ \left(\lambda \mathsf{x}. \bot, \bot\right) \right) \, \middle| \, a_0 = 0 \right\}$$

We use the notation β S to denote $\beta \cup S$. Abusing notation, we also write β (S) to refer to β or β S, e.g. if a property holds for either system. Note that for the 111 calculus, β S = β .

In addition, we continue studying the relation between m-convergence and p-convergence: In general, they are subtly different, but we show that a p-converging $\beta(S)$ -reduction can be

The official version will be available from July 9, 2018 at: http://www.dagstuhl.de/dagpub/978-3-95977-077-4 adequately simulated by an m-converging B-reduction and vice versa, where B is an extension of β , called Böhm rewrite system, which additionally contains rules of the form $t \to \bot$. This result uses the same construction used by Kennaway et al. [10] to study so-called *meaningless terms*.

In the remainder of this section we first characterise the set of lambda trees that pconverge to \perp (Section 5.1); we then establish a correspondence between pure p-convergence and m-convergence extended with rules $t \rightarrow \perp$ for lambda trees t that p-converge to \perp (Section 5.2); and finally we prove infinitary confluence and normalisation for p-convergent β S-reductions in the 001, 101, and 111 calculi (Section 5.3). For the infinitary confluence result, we make use of the correspondence between p-convergence and m-convergence.

5.1 Partiality

We begin with the characterisation of lambda trees that p-converge to \perp :

▶ **Definition 5.2.** Given an open reduction $S = (t_{\iota} \rightarrow_{p_{\iota}} t_{\iota+1})_{\iota < \alpha}$, a position p is called *volatile* in S if, for each $\beta < \alpha$, there is some $\beta \leq \gamma < \alpha$ with $p_{\gamma} \sqrt{\overline{\alpha}} \leq p \leq p_{\gamma}$. If p is volatile in S but no proper prefix of p is, then p is called *outermost-volatile* in S.

For instance, in the β -reduction in Example 4.8, $\langle 1, 0, 2 \rangle$ is volatile and $\langle 1, 0, 2 \rangle \downarrow^{\overline{a}}$ is outermost-volatile. Note that outermost-volatile positions must be non-strict, because if p is volatile, then so is $p \downarrow^{\overline{a}}$.

The presence of volatile positions characterises partiality in p-convergent reductions, which by Corollary 4.11 can be stated as follows:

▶ **Proposition 5.3.** $S: s \xrightarrow{m} t$ iff no prefix of S has volatile positions and $S: s \xrightarrow{p} t$.

Proof sketch. Let $S = (t_{\iota} \rightarrow p_{\iota} t_{\iota+1})_{\iota < \alpha}$. The "only if" direction follows from Proposition 4.9 and the fact that if $(|p_{\iota}|^{\overline{\alpha}})_{\iota < \beta}$ tends to infinity, then $S|_{\beta}$ has no volatile positions. For the "if" direction, the infinite pigeonhole principle yields that $(|p_{\iota}|^{\overline{\alpha}})_{\iota < \beta}$ tends to infinity. Using this fact, one can show that $S: s \xrightarrow{\mathbf{m}} t$.

More specifically, outermost-volatile positions pinpoint the exact location of partiality in the result of a p-converging reduction.

▶ Lemma 5.4. If p is outermost-volatile in S: $s \xrightarrow{\mathbf{p}} t$, then $p \in \mathcal{D}_{\perp}(t)$.

Proof sketch. Let $S = (t_{\iota} \to_{p_{\iota},c_{\iota}} t_{\iota+1})_{\iota < \alpha}$. Since p is volatile in S, we find for each $\beta < \alpha$ some $\beta \leq \iota < \alpha$ with $p_{\iota} \downarrow^{\overline{\alpha}} \leq p$. Hence, by Lemma 4.6, we know that $p \notin \mathcal{D}(c_{\iota})$. Consequently, by Theorem 3.4 and Proposition 3.5, we have that $p \notin \mathcal{D}(t)$. If $p = \langle \rangle$, then $p \in \mathcal{D}_{\perp}(t)$ follows immediately. If $p = q \cdot \langle 0 \rangle$, then one can use the fact that no prefix of q is volatile to show that $t(q) = \lambda$, which means that $p \in \mathcal{D}_{\perp}(t)$. The argument for the cases $p = q \cdot \langle 1 \rangle$ and $p = q \cdot \langle 2 \rangle$ is analogous.

This characterisation of partiality in terms of volatile positions motivates the following notions of destructiveness and fragility:

▶ **Definition 5.5.** A reduction S is called *destructive* if it is p-continuous and $\langle \rangle$ is volatile in S. A lambda tree $t \in \mathcal{T}_{\perp}^{\overline{a}}$ is called *fragile* if there is a destructive β -reduction starting from t. The set of all fragile *total* lambda trees is denoted $\mathcal{F}^{\overline{a}}$.

Note that fragility is defined in terms of destructive β -reductions. However, one can show that a destructive β -reduction exists iff a destructive β S-reduction exists.

The following proposition explains why destructive reductions have deserved their name:

▶ **Proposition 5.6.** An open reduction is destructive iff it p-converges to \bot .

Proof sketch. The "only if" direction follows from Lemma 5.4; the converse direction can be shown using the characterisation of the limit inferior (Theorem 3.4, Proposition 3.5).

For example, the β -reduction $\Omega \to \Omega \to \dots$ (cf. Example 4.8) p-converges to \perp and is thus destructive. As a corollary from the above proposition, we obtain that every fragile lambda tree – such as Ω – can be contracted to \perp by an open p-convergent reduction.

5.2 Correspondence

To compare m- and p-converging reductions, we employ Böhm rewrite systems and the underlying notion of \perp -instantiation from Kennaway et al.'s work on meaningless terms [10].

▶ **Definition 5.7.** Let $\mathcal{U} \subseteq \mathcal{T}^{\infty}$ and $t \in \mathcal{T}^{\infty}_{\perp}$. A lambda tree $s \in \mathcal{T}^{\infty}$ is called a ⊥-instance of t w.r.t. \mathcal{U} if s is obtained from t by inserting elements of \mathcal{U} into t at each position $p \in \mathcal{D}_{\perp}(t)$, i.e. s(p) = t(p) for all $p \in \mathcal{D}(t)$ and $s|_p \in \mathcal{U}$ for all $p \in \mathcal{D}_{\perp}(t)$. The set of lambda trees that have a ⊥-instance w.r.t. \mathcal{U} that is in \mathcal{U} itself is denoted \mathcal{U}_{\perp} . In other words, $t \in \mathcal{U}_{\perp}$ iff there is a lambda trees $s \in \mathcal{U}$ such that s is obtained from t by replacing occurrences of \perp in t by lambda trees from \mathcal{U} .

In particular, we will use the above construction with the set of fragile total lambda trees $\mathcal{F}^{\overline{a}}$, which gives us the set $\mathcal{F}^{\overline{a}}_{\perp}$.

Finally, we give the construction of Böhm rewrite systems.

▶ Definition 5.8. For each set $\mathcal{U} \subseteq \mathcal{T}^{\overline{n}}$, we define the following two rewrite systems:

$$\mathbb{I}(\mathcal{U}) = \{(t, \bot) \mid t \in \mathcal{U}_{\bot} \setminus \{\bot\}\}, \qquad \qquad \mathbb{B}(\mathcal{U}) = \beta \cup \mathbb{I}(\mathcal{U})$$

If \mathcal{U} is clear from the context, we instead use the notation \bot and \mathbb{B} , respectively.

In particular, we consider the Böhm rewrite system w.r.t. fragile total lambda trees, denoted by $\mathbb{B}(\mathcal{F}^{\overline{a}})$. We start with one direction of the correspondence between **p**-converging $\beta(\mathbb{S})$ reductions and **m**-converging $\mathbb{B}(\mathcal{F}^{\overline{a}})$ -reductions:

▶ Theorem 5.9. If $s \xrightarrow{\mathbf{p}}_{\mathfrak{B}S} t$, then $s \xrightarrow{\mathbf{m}}_{\mathfrak{B}} t$, where $\mathbb{B} = \mathbb{B}(\mathcal{F}^{\overline{a}})$.

Proof sketch. Given $S: s \xrightarrow{\mathsf{Pu}}_{\beta \mathbb{S}} t$, we construct a B-reduction T from S that also p-converges to t but that has no volatile positions in any of its open prefixes. Thus, according to Proposition 5.3, $T: s \xrightarrow{\mathsf{m}}_{\mathcal{B}} t$. The construction of T removes steps in S that take place at or below any outermost-volatile position of some prefix of S and replaces them by a single \mathbb{I} -step. Such a \mathbb{I} -step can be performed since a fragile lambda tree must be responsible for an outermost-volatile position. Moreover, \mathbb{S} -steps in S are \mathbb{I} -steps in T since $\mathbb{S} \subseteq \mathbb{I}(\mathcal{F}^{\overline{a}})$. Lemma 5.4 can then be used to show that the resulting B-reduction T p-converges to t.

The converse direction of Theorem 5.9 does not hold in general. The problem is that \mathbb{I} -steps can be more selective in which fragile lambda subtree to contract to \perp compared to p-convergent reductions with volatile positions. If p is a volatile position, then so is $p\downarrow^{\overline{a}}$. Consequently, volatile positions and thus ' \perp 's in the result of a p-converging reduction are propagated upwards through strict positions. For example, let $a_0 = 0$, and $t = \lambda y.\Omega$. Since Ω is fragile, we have the reduction $t \to_{\mathbb{I}} \lambda y. \bot$. On the other hand, via p-convergent β -reductions, t only reduces to itself and \perp . This phenomenon, however, does not occur if we restrict ourselves to the strictness signature 111 or if we only consider \mathbb{I} -normal forms. Indeed, in the above example, $\lambda y. \perp$ is not a \mathbb{I} -normal form and can be contracted to \perp with a \mathbb{I} -step.

The official version will be available from July 9, 2018 at: http://www.dagstuhl.de/dagpub/978-3-95977-077-4

P. Bahr

▶ Theorem 5.10. Let $\mathbb{B} = \mathbb{B}(\mathcal{F}^{\overline{a}})$ and $s \xrightarrow{\mathbb{m}}_{\mathbb{B}} t$ such that s is total. Then $s \xrightarrow{\mathbb{P}}_{\mathbb{B}} t$ if $\overline{a} = 111$ or t is a \mathbb{I} -normal form.

Proof sketch. The reduction $s \xrightarrow{\mathbf{m}}_{\mathcal{B}} t$ can be factored into $S: s \xrightarrow{\mathbf{m}}_{\mathcal{B}} s'$ and $T: s' \xrightarrow{\mathbf{m}}_{\mathcal{A}} t$ (by the same proof as Lemma 27 of Kennaway et al. [10]). Moreover, we may assume w.l.o.g. that T contracts disjoint \bot -redexes in s' (using an argument similar to Lemma 7.2.4 of Ketema [11]). By Proposition 4.9, we have that $S: s \xrightarrow{\mathbf{p}}_{\mathcal{B}} s'$ and that $T: s' \xrightarrow{\mathbf{p}}_{\mathcal{A}} t$. For each step $u \to_{\bot,p} v$ in T we find a reduction $T_p: u \xrightarrow{\mathbf{p}}_{\mathcal{B}} v'$ in which p is volatile since $u|_p$ must be fragile. Given that $\overline{a} = 111$ or that t is a \bot -normal form, we can show that p is in fact outermost-volatile in T_p . Hence, the equality v = v' follows from Lemma 5.4. Therefore, we may replace each step $u \to_{\bot,p} v$ in T by T_p , which yields a reduction $s' \xrightarrow{\mathbf{p}}_{\mathcal{B}} t$.

That is, in general we get one direction of the correspondence – namely from metric to partial order reduction – only for reductions to normal forms. However, this does not matter that much as p-converging $\beta(S)$ -reductions (an thus also m-converging $\mathbb{B}(\mathcal{F}^{\overline{a}})$ -reductions) are normalising as we show below.

5.3 Infinitary Normalisation and Confluence

We begin by recalling the notion of active lambda trees [10], which we use to establish infinitary normalisation and as an alternative characterisation of fragile lambda trees (in the 001, 101, and 111 calculi).

▶ **Definition 5.11.** A lambda tree t is called *stable* if no lambda tree t' with $t \to_{\beta}^{*} t'$ has a β -redex occurrence at \overline{a} -depth 0; t is called *active* if no lambda tree t' with $t \to_{\beta}^{*} t'$ is stable. The set of all active *total* lambda trees is denoted by $\mathcal{A}^{\overline{a}}$.

To construct normalising p-convergent reductions, we follow the idea of Kennaway et al. [10]: We contract all subtrees of the initial lambda tree into stable form. The only way to achieve this for active subtrees is to annihilate them by a destructive reduction. The basis for that strategy is the following observation:

▶ Lemma 5.12. Every active lambda tree is fragile.

Proof. If t_0 is active, we find a reduction $t_0 \rightarrow^*_{\beta} t'_0$ to a β -redex at \overline{a} -depth 0. By contracting this redex we get a lambda tree t_1 that is active, too. By repeating this argument we obtain a destructive reduction $t_0 \rightarrow^*_{\beta} t'_0 \rightarrow_{\beta} t_1 \rightarrow^*_{\beta} t'_1 \rightarrow_{\beta} \dots$

The following normalisation result then follows straightforwardly:

▶ Theorem 5.13. For each $s \in \mathcal{T}_{\perp}^{\overline{a}}$, there is a normalising reduction $s \stackrel{\mathsf{p}}{\xrightarrow{}}_{\beta(\mathbb{S})} t$.

Proof sketch. Similar to Theorem 1 of Kennaway et al. [10]: an active subtree at position p is by Lemma 5.12 also fragile. Hence, there is a β -reduction in which a prefix of p is outermost-volatile. By Lemma 5.4, such a reduction annihilates the active subtree at p. This yields a reduction $s \xrightarrow{\mathbf{p}} t$ to β -normal form t, which can be extended by a reduction $t \xrightarrow{\mathbf{p}} u$ to a β S-normal form u.

From the above we immediately obtain the corresponding result for m-convergence:

▶ Theorem 5.14. For each $s \in \mathcal{T}_{\perp}^{\overline{a}}$ there is a normalising reduction $s \xrightarrow{\mathsf{m}}_{\mathbb{B}(\mathcal{F}^{\overline{a}})} t$.

Proof. By Theorem 5.13 and 5.9, as β S-normal forms are also $\mathbb{B}(\mathcal{F}^{\overline{a}})$ -normal forms.

FSCD 2018

Consequently, we can derive the following correspondence result.

► Corollary 5.15. For each $s \in \mathcal{T}^{\overline{a}}$ with $s \xrightarrow{\mathfrak{m}}_{\mathbb{B}(\mathcal{F}^{\overline{a}})} t$, there is a reduction $t \xrightarrow{\mathfrak{m}}_{\mathbb{B}(\mathcal{F}^{\overline{a}})} t'$ such that $s \xrightarrow{\mathfrak{p}}_{\mathfrak{m}} t'$.

Proof. According to Theorem 5.14, there is a normalising reduction $t \xrightarrow{\mathbf{m}}_{\mathbb{B}(\mathcal{F}^{\overline{a}})} t'$. Then a reduction $s \xrightarrow{\mathbf{p}}_{\mathbb{B}} t'$ exists by Theorem 5.10.

A shortcoming of this correspondence property and the correspondence properties established in Section 5.2 is that they consider m-convergence in the system $\mathbb{B}(\mathcal{F}^{\overline{a}})$, which is unsatisfactory since $\mathcal{F}^{\overline{a}}$ is defined using p-convergence. A more appropriate choice would be the set $\mathcal{A}^{\overline{a}}$ of active terms, which is defined in terms of finitary reduction only. To obtain a correspondence in terms of $\mathcal{A}^{\overline{a}}$, we will show that $\mathcal{F}^{\overline{a}} = \mathcal{A}^{\overline{a}}$ for strictness signatures 001, 101, and 111. To prove this equality of fragility and activeness, we need the following key lemma, which can be proved using descendants and complete developments.

▶ Lemma 5.16 (Infinitary Strip Lemma). Given $S: s \xrightarrow{\mathsf{P}_{\mathsf{W}}} t_1$ and $T: s \rightarrow_{\beta S}^* t_2$, there are reductions $S': t_1 \xrightarrow{\mathsf{P}_{\mathsf{W}}} t_3$ and $T': t_2 \xrightarrow{\mathsf{P}_{\mathsf{W}}} t_3$, provided $\overline{a} \in \{001, 101, 111\}$.

Recall that $\beta S = \beta$ for $\overline{a} = 111$, i.e. the infinitary strip lemma holds for pure β -reduction in the 111 calculus; but it does not hold for 001 and 101 as Example 5.1 demonstrates. Hence, the need for S-rules. By contrast, in the metric calculi of Kennaway et al. [9] the infinitary strip lemma does not hold for any \overline{a} . In order to obtain the infinitary strip lemma and confluence, Kennaway et al. extended β -reduction to Böhm reduction.

We use the Infinitary Strip Lemma to show that p-convergent reductions to \perp can be compressed to length at most ω .

▶ Lemma 5.17. If $\overline{a} \in \{001, 101, 111\}$ and $S: t \xrightarrow{P_{*}}_{\beta S} \bot$, then there is a reduction $T: t \xrightarrow{P_{*}}_{\beta S} \bot$ of length $\leq \omega$. If t is total, then T is a β -reduction of length ω .

Proof sketch. If $|S| \leq \omega$, we are done. Otherwise, we can construct a finite reduction $t \to_{\beta\mathbb{S}}^* t'$ with at least one contraction at \overline{a} -depth 0 either using a finite approximation property of p-convergence (in case S contracts β -redex at \overline{a} -depth 0) or by an induction argument (in case S contracts \mathbb{S} -redex at root position). By Lemma 5.16, there is a reduction $S': t' \xrightarrow{\mathbb{P}}_{\beta\mathbb{S}} \bot$. Thus, we can repeat the argument for S'. Iterating this argument yields either a reduction $t \to_{\beta\mathbb{S}}^* \bot$ or a reduction $t \xrightarrow{\mathbb{P}}_{\beta\mathbb{S}} s'$ of length ω with infinitely many contractions at \overline{a} -depth 0, and thus $s' = \bot$. If s is total, then T cannot be finite, as finite $\beta\mathbb{S}$ -reductions preserve totality. Hence, no step in T can be an \mathbb{S} -step.

▶ Lemma 5.18. If $\overline{a} \in \{001, 101, 111\}$, a total lambda tree is active iff it is fragile.

Proof. The "only if" direction follows from Lemma 5.12. For the converse direction let t be total and fragile, and let $t \to_{\beta}^{*} t_1$. Since t is fragile, there is a reduction $t \stackrel{\mathbf{p}_{\to}}{}_{\beta \mathbb{S}} \perp$ according to Proposition 5.6. Hence, by Lemma 5.16, there is a reduction $T: t_1 \stackrel{\mathbf{p}_{\to}}{}_{\beta \mathbb{S}} \perp$, which we can assume, according to Lemma 5.17, to be a β -reduction of length ω . Since T is, by Proposition 5.6, destructive, there is a proper prefix $T': t_1 \stackrel{\mathbf{p}_{\to}}{}_{\beta} t_2$ of T such that t_2 has a redex occurrence at \overline{a} -depth 0. Because T is of length ω , T' is finite i.e. $T': t_1 \to_{\beta}^{*} t_2$.

The above lemma allows us to derive confluence w.r.t. p-convergent reductions from the confluence results w.r.t. m-convergence of Kennaway et al. [9]:

▶ **Theorem 5.19** (infinitary confluence). Given $\overline{a} \in \{001, 101, 111\}$, we have that $s \xrightarrow{\mathbf{p}}_{\mathsf{M}} \mathfrak{g} t_1$ and $s \xrightarrow{\mathbf{p}}_{\mathsf{M}} \mathfrak{g} t_2$ implies that $t_1 \xrightarrow{\mathbf{p}}_{\mathsf{M}} \mathfrak{g} t_1$ and $t_2 \xrightarrow{\mathbf{p}}_{\mathsf{M}} \mathfrak{g} t_1$.

The official version will be available from July 9, 2018 at: http://www.dagstuhl.de/dagpub/978-3-95977-077-4 **Proof.** According to Theorem 5.13, we can extend the existing reductions by normalising reductions $t_1 \xrightarrow{\mathbf{p}}_{\mathcal{P} \mathcal{P} \mathcal{S}} t'_1$ and $t_2 \xrightarrow{\mathbf{p}}_{\mathcal{P} \mathcal{P} \mathcal{S}} t'_2$. By Theorem 5.9 and Lemma 5.18, the resulting normalising reductions $s \xrightarrow{\mathbf{p}}_{\mathcal{P} \mathcal{P} \mathcal{S}} t'_1$ and $s \xrightarrow{\mathbf{p}}_{\mathcal{P} \mathcal{P} \mathcal{S}} t'_2$ are also m-convergent $\mathbb{B}(\mathcal{A}^{\overline{a}})$ -reductions. Kennaway et al. [9] have shown that such reductions are confluent. Hence, $t'_1 = t'_2$ (as βS -normal forms are $\mathbb{B}(\mathcal{A}^{\overline{a}})$ -normal forms too).

Together with the earlier normalisation result, this means that the 001, 101, and 111 calculi have unique normal forms w.r.t. $\overset{p}{\longrightarrow}_{\beta S}$. By the correspondence results between the metric and the partial order calculi, these normal forms are the same as the unique normal forms w.r.t. $\overset{m}{\longrightarrow}_{\mathbb{B}(\mathcal{A}^{\overline{a}})}$ [9], which in turn correspond to Böhm Trees, Levy-Longo Trees, and Berarducci Trees, respectively.

6 Related Work

The use of ideal completion in lambda calculus to construct infinite terms has a long history (see e.g. Ketema [11] for an overview), in particular in the form of constructing infinite normal forms such as Böhm Trees. In that line of work, the ideal completion is typically based on the fully monotone partial order \leq_{\perp} generated by $\perp \leq_{\perp} M$ for any term M. Different kinds of infinite normal forms are then obtained by modulating the set of rules that are used to generate the normal forms. In this paper, we instead modulated the partial order and we have constructed full infinitary calculi in the style of Kennaway et al. [9]. Blom's abstract theory of infinite normal forms and infinitary rewriting based on ideal completion [8] has been crucial for developing our infinitary calculi.

In previous work, we have compared infinitary rewriting based on partial orders vs. metric spaces in a first-order setting [3, 4]. However, in that work we have only considered fully non-strict convergence, whereas we consider varying modes of strictness in the present paper.

Blom's work [16] on preservation calculi is similar to our ideal completion calculi. Blom also considers different calculi indexed by strictness signatures and relates them to the corresponding metric calculi. However, he uses the same partial order \leq_{\perp}^{111} for all calculi; the different calculi vary in the notion of reduction contexts they use. Blom's reduction contexts are the same as our reduction contexts, and his Ω -rules are more general variants of our S-rules. However, his approach of using a single partial order has some caveats:

Firstly, there is no corresponding weak notion of preservation sequences that corresponds to weak m-convergence. Secondly, the partially ordered set $(\mathcal{T}_{\perp}^{\overline{a}}, \trianglelefteq_{\perp}^{111})$ is only a complete semilattice for $\overline{a} = 111$; otherwise it is not even a cpo and limit inferiors do not always exist. For example, let t be an \overline{a} -unguarded lambda tree (i.e. $t \notin \mathcal{T}_{\perp}^{\overline{a}}$), and for each $i < \omega$ let t_i be the restriction of t to positions of depth < i, which means that $t_i \in \mathcal{T}_{\perp}^{\overline{a}}$. Then liminf_{$i\to\omega$} t_i w.r.t. $\trianglelefteq_{\perp}^{111}$ is t itself and thus not in $\mathcal{T}_{\perp}^{\overline{a}}$ even though all t_i are. This does not cause a problem, if one only considers reduction contexts of p-continuous reductions, though.

For the comparison of his preservation calculi with the metric calculi, Blom uses a notion of 0-active terms, which is different from the notion of active terms as used here and by Kennaway et al. [9, 10] (under the names 0-activeness resp. *abc*-activeness). Blom defines that a lambda tree is 0-active iff there is a destructive reduction of length ω starting from it. 0-activeness is demonstrably different from activeness for any strictness signature with $a_2 = 0$ as Example 5.1 shows. But 0-activeness and activeness do coincide for 001, 101, and 111 as we have shown with the combination of Lemma 5.17 and Lemma 5.18.

— References

- André Arnold and Maurice Nivat. The metric space of infinite trees. algebraic and topological properties. *Fundamenta Informaticae*, 3(4):445–476, 1980.
- 2 Patrick Bahr. Abstract models of transfinite reductions. In Christopher Lynch, editor, Proceedings of the 21st International Conference on Rewriting Techniques and Applications, RTA 2010, July 11-13, 2010, Edinburgh, Scottland, UK, volume 6 of LIPIcs, pages 49– 66. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010. doi:10.4230/LIPIcs.RTA. 2010.49.
- 3 Patrick Bahr. Partial order infinitary term rewriting and böhm trees. In Christopher Lynch, editor, Proceedings of the 21st International Conference on Rewriting Techniques and Applications, RTA 2010, July 11-13, 2010, Edinburgh, Scottland, UK, volume 6 of LIPIcs, pages 67–84. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010. doi:10. 4230/LIPIcs.RTA.2010.67.
- Patrick Bahr. Infinitary term graph rewriting is simple, sound and complete. In Ashish Tiwari, editor, 23rd International Conference on Rewriting Techniques and Applications (RTA'12), RTA 2012, May 28 June 2, 2012, Nagaya, Japan, volume 15 of LIPIcs, pages 69–84. Schloss Dagstuhl Leibniz-Zentrum fuer Informatik, 2012. doi:10.4230/LIPIcs. RTA.2012.69.
- 5 Patrick Bahr. Ideal completions of the lambda calculus. Companion report, 2018. URL: https://arxiv.org/abs/1805.06736.
- 6 Henk P Barendregt. The Lambda Calculus: Its Syntax and Semantics, volume 103 of Studies in Logic and the Foundations of Mathemantics. Elsevier Science, revised edition, 1984.
- 7 Alessandro Berarducci. Infinite A-calculus and non-sensible models. In A Ursini and P Aglianó, editors, *Logic and algebra*, number 180 in Lecture Notes in Pure and Applied Mathematics, pages 339–378. CRC Press, 1996.
- 8 Stefan Blom. Term Graph Rewriting-Syntax and Semantics. PhD thesis, Vrije Universiteit te Amsterdam, 2001.
- 9 Richard Kennaway, Jan Willem Klop, M R Sleep, and Fer-Jan de Vries. Infinitary lambda calculus. *Theoretical Computer Science*, 175(1):93-125, 1997. doi:DOI:10.1016/ S0304-3975(96)00171-5.
- 10 Richard Kennaway, Vincent van Oostrom, and Fer-Jan de Vries. Meaningless terms in rewriting. *Journal of Functional and Logic Programming*, 1999(1):1–35, 1999.
- 11 Jeroen Ketema. *Böhm-Like Trees for Rewriting*. PhD thesis, Vrije Universiteit Amsterdam, 2006. URL: http://dare.ubvu.vu.nl/handle/1871/9203.
- 12 Jean-Jacques Lévy. An algebraic interpretation of the *lambda beta* k-calculus; and an application of a labelled *lambda* -calculus. *Theor. Comput. Sci.*, 2(1):97–114, 1976. doi: 10.1016/0304-3975(76)90009-8.
- 13 Giuseppe Longo. Set-theoretical models of λ -calculus: theories, expansions, isomorphisms. Annals of pure and applied logic, 24(2):153–188, 1983.
- 14 Jean-Jacques Lévy. Réductions Correctes et Optimales dans le Lambda-Calcul. PhD thesis, Université Paris VII, 1978.
- 15 Mila E Majster-Cederbaum and Christel Baier. Metric completion versus ideal completion. *Theoretical Computer Science*, 170(1-2):145–171, 1996. doi:DOI:10.1016/S0304-3975(96) 80705-5.
- 16 Vincent van Oostrom, editor. Rewriting Techniques and Applications, 15th International Conference, RTA 2004, Aachen, Germany, June 3-5, 2004, Proceedings, volume 3091 of Lecture Notes in Computer Science. Springer, 2004. doi:10.1007/b98160.