Spinal Atomic Lambda-Calculus

David Sherratt¹ (✉), Willem Heijltjes², Tom Gundersen³, and Michel Parigot⁴

¹ Friedrich-Schiller-Universität Jena, Germany.

david.rhys.sherratt@uni-jena.de

² University of Bath, United Kingdom.

w.b.heijltjes@bath.ac.uk

³ Red Hat, Inc. Norway.

teg@jklm.no

⁴ Institut de Recherche en Informatique Fondamentale, CNRS, Université de Paris.
France.

parigot@irif.fr

Abstract. We present the spinal atomic λ -calculus, a typed λ -calculus with explicit sharing and atomic duplication that achieves spinal full laziness: duplicating only the direct paths between a binder and bound variables is enough for beta reduction to proceed. We show this calculus is the result of a Curry–Howard style interpretation of a deep-inference proof system, and prove that it has natural properties with respect to the λ -calculus: confluence and preservation of strong normalisation.

Keywords: Lambda-Calculus · Full laziness · Deep inference · Curry–Howard

1 Introduction

In the λ -calculus, a main source of efficiency is *sharing*: multiple use of a single subterm, commonly expressed through graph reduction [27] or explicit substitution [1]. This work, and the *atomic λ -calculus* [16] on which it builds, is an investigation into sharing as it occurs naturally in intuitionistic *deep-inference* proof theory [26]. The atomic λ -calculus arose as a Curry–Howard interpretation of a deep-inference proof system, in particular of the *distribution* rule given below left, a variant of the characteristic *medial* rule [10, 26]. In the term calculus, the corresponding *distributor* enables duplication to proceed *atomically*, on individual constructors, in the style of sharing graphs [21]. As a consequence, the natural reduction strategy in the atomic λ -calculus is *fully lazy* [27, 4]: it duplicates only the minimal part of a term, the *skeleton*, that can be obtained by lifting out subterms as explicit substitutions. (While duplication is atomic *locally*, a duplicated abstraction does not form a redex until also its bound variables have been duplicated; hence duplication becomes fully lazy *globally*.)

This work was supported by EPSRC Project EP/R029121/1 *Typed Lambda-Calculi with Sharing and Unsharing* and ANR project 15-CE25-0014 *The Fine Structure of Formal Proof Systems and their Computational Interpretations (FISP)*

$$\text{Distribution: } \frac{A \rightarrow (B \wedge C)}{(A \rightarrow B) \wedge (A \rightarrow C)}^d \quad \text{Switch: } \frac{(A \rightarrow B) \wedge C}{A \rightarrow (B \wedge C)}^s$$

We investigate the computational interpretation of another characteristic deep-inference proof rule: the *switch* rule above right [26].⁵ Our result is the *spinal atomic λ -calculus*, a λ -calculus with a refined form of full laziness, *spine duplication*. In the terminology of [4], this strategy duplicates only the *spine* of an abstraction: the paths to its bound variables in the syntax tree of the term.⁶

We illustrate these notions in Figure 1, for the example $\lambda x.\lambda y.((\lambda z.z)y)x$. The *scope* of the abstraction λx is the entire subterm, $\lambda y.((\lambda z.z)y)x$ (which may or may not be taken to include λx itself). Note that with explicit substitution, the scope may grow or shrink by lifting explicit substitutions in or out. The *skeleton* is the term $\lambda x.\lambda y.(wy)x$ where the subterm $\lambda z.z$ is lifted out as an (explicit) substitution $[\lambda z.z/w]$. The *spine* of a term, indicated in the second image, cannot naturally be expressed with explicit substitution, though one can get an impression with *capturing* substitutions: it would be $\lambda x.\lambda y.wx$, with the subterm $(\lambda z.z)y$ extracted by a capturing substitution $[(\lambda z.z)y/w]$. Observe that the skeleton can be described as the *iterated spine*: it is the smallest subgraph of the syntax tree closed under taking the spine of each abstraction, i.e. that contains the spine of every abstraction it contains.

These notions give rise to four natural duplication regimes. For a shared abstraction to become available as the function in a β -redex: *laziness* duplicates its *scope* [22]; *Full laziness* duplicates its *skeleton* [27]; *Spinal full laziness* duplicates its *spine* [8]; *optimal reduction* duplicates only the abstraction λx and its bound variables x [21, 3].⁷

While each of these duplication strategies has been expressed in graphs and labelled calculi, the atomic λ -calculus is the first term calculus with Curry–Howard corresponding proof system to naturally describe full laziness. Likewise, the spinal atomic λ -calculus presented here is the first term calculus with Curry–Howard corresponding proof system to naturally describe spinal full laziness.

Switch and Spine. One way to describe the skeleton or the spine of an abstraction within a λ -term is through explicit end-of-scope markers, as explored by Berkling and Fehr [7], and more recently by Hendriks and Van Oostrom [18]. We use their *adbm*al (λ) to illustrate the idea: the constructor $\lambda x.N$ indicates that the subterm N does not contain occurrences of x (or that any that do occur are

⁵ The switch rule is an intuitionistic variant of *weak* or *linear distributivity* [12] for multiplicative linear logic.

⁶ There is a clash of (existing) terminology: the *spine of an abstraction*, as we use here, is a different notion from the *spine of a λ -term*, which is the path from the root to the leftmost variable, as used e.g. in head reduction and abstract machines.

⁷ Interestingly, Balabonski [5] shows that for *weak* reduction (where one does not reduce under an abstraction) full laziness and spinal full laziness are both optimal (in the number of beta-steps required to reach a normal form).

Fig. 1: Balanced and unbalanced typing derivations for $\lambda x.\lambda y.((\lambda z.z)y)x$, with corresponding graphical representations of the term. The variable x has type A and y, z type $A \rightarrow B$, shortened to B^A . The left derivation isolates the skeleton of λx , and the right derivation its spine, both by the subderivations in braces.

not available to a binder λx outside $\lambda x.N$). The scope of an abstraction thus becomes explicitly indicated in the term. This opens up a distinction between *balanced* and *unbalanced* scopes: whether scopes must be properly nested, or not; for example, in $\lambda x.\lambda y.N$, a subterm $\lambda y.\lambda x.M$ is balanced, but $\lambda x.\lambda y.M$ is not. With balanced scope, one can indicate the skeleton of an abstraction; with unbalanced scope (which Hendriks and Van Oostrom dismiss) one can indicate the spine. We do so for our example term $\lambda x.\lambda y.((\lambda z.z)y)x$ below.

Balanced scope/skeleton: $\lambda x.\lambda y.(\lambda y.(\lambda x.\lambda z.z)y)(\lambda y.x)$

Unbalanced scope/spine: $\lambda x.\lambda y.(\lambda x.(\lambda y.\lambda z.z)y)(\lambda y.x)$

A closely related approach is *director strings*, introduced by Kennaway and Sleep [19] for combinator reduction and generalized to any reduction strategy by Fernández, Mackie, and Sinot in [13]. The idea is to use nameless abstractions identified by their nesting (as with De Bruijn indices), and make the paths to bound variables explicit by annotating each constructor with a string of *directors*, that outline the paths. The primary aim of these approaches is to eliminate α -conversion and to streamline substitution. Consequently, while they can *identify* the spine, they do not readily isolate it for duplication.

The present work starts from our observation that the *switch* rule of open deduction functions as a proof-theoretic end-of-scope construction (see [25] for details). However, it does so in a *structural* way: it forces a deconstruction of a proof into readily duplicable parts, which together may form the spine of an abstraction. The derivations in Figure 1 demonstrate this, as we will now explain—see the next section for how they are formally constructed.

The abstraction λx corresponds in the proof system to the implication $A \rightarrow$, explicitly scoping over its right-hand side. On the left, with the *abstraction* rule (λ), scopes must be balanced, and the proof system may identify the *skeleton*; here, that of λx as the largest blue box. Decomposing the abstraction (λ) into *axiom* (a) and *switch* (s), on the right the proof system may express unbalanced

scope. It does so by separating the scope of an abstraction into multiple parts; here, that of λx is captured as the two top-level red boxes. Each box is ready to be duplicated; in this way, one may duplicate the spine of an abstraction only.

These two derivations correspond to terms in our calculus. The subterms not part of the skeleton (i.e. $\lambda z.z$) remain shared and we are able to duplicate the skeleton alone. This is also possible in [16]. In our calculus we are also able to duplicate just the spine by using a *distributor*. We require this construct as otherwise we break the binding of the y -abstraction. The distributor manages and maintains these bindings. The y -abstraction in the spine ($y\langle a \rangle$) is a *phantom-abstraction*, because it is not real and we cannot perform β -reduction on it. However, it may become real during reduction. It can be seen as a placeholder for the abstraction. The variables in the *cover* (a) represent subterms that both remain shared and are found in the distributor.

$$\begin{array}{ll} \text{Skeleton:} & \underline{\lambda x. \lambda y. (a y)} x [a \leftarrow \lambda z. z] \\ \text{Spine:} & \underline{\lambda x. y\langle a \rangle. (a)} x [y\langle a \rangle \mid \lambda y. [a \leftarrow (\lambda z. z) y]] \end{array}$$

Our investigation is then focused on the interaction of switch and distribution (later observed in the rewrite rule l_5). The use of the distribution rule allows us to perform duplication atomically, and thus provides a natural strategy for spinal full laziness. In Figure 1 on the right, this means duplicating the two top-level red boxes can be done independently from duplicating the yellow box.

2 Typing a λ -calculus in open deduction

We work in *open deduction* [15], a formalism of deep-inference proof theory, using the following proof system for (conjunction–implication) intuitionistic logic. A *derivation* from a *premise* formula X to a *conclusion* formula Z is constructed inductively as in Figure 2a, with from left to right: a propositional atom a , where $X = Z = a$; *horizontal composition* with a connective $\rightarrow$, where $X = Y \rightarrow X_2$ and $Z = Y \rightarrow Z_2$; *horizontal composition* with a connective $\wedge$, where $X = X_1 \wedge X_2$ and $Z = Z_1 \wedge Z_2$; and *rule composition*, where r is an inference rule (Figure 2b) from Y_1 to Y_2 . The boxes serve as parentheses (since derivations extend in two dimensions) and may be omitted. Derivations are considered up to associativity of rule composition. One may consider formulas as derivations that omit rule composition. We work modulo associativity, symmetry, and unitality of conjunction, justifying the n -ary contraction, and may omit $\top$ from the axiom rule. A 0-ary contraction, with conclusion $\top$, is a *weakening*. Figure 2b: the abstraction rule (λ) is derived from axiom and switch. *Vertical composition* of a derivation from X to Y and one from Y to Z , depicted by a dashed line, is a defined operation, given in Figure 2c, where $* \in \{\wedge, \rightarrow\}$.

2.1 The Sharing Calculus

Our starting point is the *sharing calculus* (λ^S), a calculus with an explicit sharing construct, similar to explicit substitution.

Fig. 2: Intuitionistic proof system in open deduction

Definition 1. The **pre-terms** r, s, t, u and **sharings** $[\Gamma]$ of the Λ^S are defined by:

$$s, t ::= x \mid \lambda x. t \mid s t \mid t[\Gamma] \quad [\Gamma] ::= [x_1, \dots, x_n \leftarrow s]$$

with from left to right: a **variable**; an **abstraction**, where x occurs free in t and becomes bound; an **application**, where s and t use distinct variable names; and a **closure**; in $t[\vec{x} \leftarrow s]$ the variables in the vector $\vec{x} = x_1, \dots, x_n$ all occur in t and become bound, and s and t use distinct variable names. **Terms** are pre-terms modulo **permutation** equivalence ($\sim$):

$$t[\vec{x} \leftarrow s][\vec{y} \leftarrow r] \sim t[\vec{y} \leftarrow r][\vec{x} \leftarrow s] \quad (\{\vec{y}\} \cap (s)_{fv} = \{\})$$

A term is in **sharing normal form** if all sharings occur as $[\vec{x} \leftarrow x]$ either at the top level or directly under a binding abstraction, as $\lambda x. t[\vec{x} \leftarrow x]$.

Note that variables are *linear*: variables occur at most once, and bound variables must occur. A vector $\vec{x}$ has length $|\vec{x}|$ and consist of the variables $x_1, \dots, x_{|\vec{x}|}$. An **environment** is a sequence of sharings $[\Gamma] = [\Gamma_1] \dots [\Gamma_n]$. Substitution is written $\{t/x\}$, and $\{t_1/x_1\} \dots \{t_n/x_n\}$ may be abbreviated to $\{t_i/x_i\}_{i \in [n]}$.

Definition 2. The **interpretation** $\llbracket - \rrbracket : \Lambda \rightarrow \Lambda^S$ is defined below.

$$\llbracket x \rrbracket = x \quad \llbracket \lambda x. t \rrbracket = \lambda x. \llbracket t \rrbracket \quad \llbracket s t \rrbracket = \llbracket s \rrbracket \llbracket t \rrbracket \quad \llbracket t[\vec{x} \leftarrow s] \rrbracket = \llbracket t \rrbracket \{ \llbracket s \rrbracket / x_i \}_{i \in [n]}$$

The **translation** $\langle N \rangle$ of a λ -term N is the unique sharing-normal term t such that $N = \llbracket t \rrbracket$. A term t will be typed by a derivation with restricted types,

Basic Types: $A, B, C := a \mid A \rightarrow B$

 Context Types: $\Gamma, \Delta, \Omega := A \mid \top \mid \Gamma \wedge \Delta$

$$\begin{array}{c}
 x : A^x \quad ts : \frac{\frac{\Gamma \parallel_t A \rightarrow B}{\wedge} \quad \frac{\Delta \parallel_s A}{\wedge}}{B} \circledast \quad \lambda x.t : \frac{\Gamma}{A \rightarrow \frac{\Gamma \wedge A^x}{\parallel_t B}} \lambda \quad t[\tilde{x} \leftarrow s] : \frac{\Gamma \wedge \frac{\Delta \parallel_s A}{\wedge} \quad \frac{\Gamma \wedge A^x}{\parallel_t B}}{\Gamma \wedge (A \wedge \dots \wedge A)^{\tilde{x}} \parallel_t B}
 \end{array}$$

 Fig. 3: Typing system for Λ^S

as shown below, where the *context type* $\Gamma = A_1 \wedge \dots \wedge A_n$ will have an A_i for each free variable x_i of t . We connect free variables to their premises by writing A^x and $\Gamma^{\tilde{x}}$. The Λ^S is then typed as in Figure 3.

3 The Spinal Atomic λ -Calculus

We now formally introduce the syntax of the spinal atomic λ -calculus (Λ_a^S), by extending the definition of the sharing calculus in Definition 1 with a *distributor* construct that allows for atomic duplication of terms.

Definition 3 (Pre-Terms). *The pre-terms r, s, t , closures $[\Gamma]$, and environments $\overline{[\Gamma]}$ of the Λ_a^S are defined by:*

$$\begin{array}{lcl}
 t ::= x \mid st \mid x(\tilde{y}).t \mid t[\Gamma] & \overline{[\Gamma]} ::= [\Gamma] \mid \overline{[\Gamma]}[\Gamma] \\
 [\Gamma] ::= [\tilde{x} \leftarrow t] \mid [\tilde{x} \mid y(\tilde{z}) \overline{[\Gamma]}]
 \end{array}$$

Our generalized abstraction $x(\tilde{y}).t$ is a **phantom-abstraction**, where x a **phantom-variable** and the **cover** $\tilde{y}$ will be a subset of the free variables of t . It can be thought of as a “delayed” abstraction: x is a binder, but possibly not in t itself, and instead in the terms substituted for the variables $\tilde{y}$; in other words, x is a *capturing* binder for substitution into $\tilde{y}$. We define standard λ -abstraction as the special case $\lambda x.t \equiv x(x).t$, and generally, when we refer to $x(\tilde{y})$ as a phantom-abstraction (rather than an abstraction) we assume $\tilde{y} \neq x$. The **distributor** $u[\tilde{x} \mid y(\tilde{z}) \overline{[\Gamma]}]$ binds the phantom-variables $\tilde{x}$ in u , while its environment $\overline{[\Gamma]}$ will bind the variables in their covers; intuitively, it represents a set of explicit substitutions in which the variables $\tilde{x}$ are expected to be captured.

The distributor is introduced when we wish to duplicate an abstraction, as depicted in Figure 4a. The sharing node ($\circ$) duplicates the abstraction node, creating a distributor (depicted as the sharing and unsharing node ($\bullet$), together with the bindings of the phantom-variables (depicted with a dashed line). The variables captured by the environment are the variables connected to sharing nodes linked with a dotted line. Notice one sharing node can be linked with multiple unsharing nodes, and vice versa. Duplication of applications also duplicates

Fig. 4: Graphical illustration of the distributor

the dotted line (Figure 4b), but these can be removed later if the term does not contain the variable bound to the unsharing (Figure 4c). These subterms are those which are not part of the spine. Eventually, we will reach a state where the only sharing node connected to the unsharing node is the one that shared the variable bound to the unsharing, allowing us to eliminate the distributor (Figure 4d). The purpose of the dotted line is similar to the brackets of optimal reduction graphs [21, 24], to supervise which sharing and unsharing match.

Terms are then pre-terms with sensible and correct bindings. To define terms, we first define *free* and *bound* variables and phantom variables; variables are bound by abstractions (not phantoms) and by sharings, while phantom-variables are bound by distributors.

Definition 4 (Free and Bound Variables). The *free variables* $(-)_f$ and *bound variables* $(-)_{bv}$ of a pre-term t are defined as follows

$$\begin{aligned}
 (x)_{fv} &= \{x\} & (x)_{bv} &= \{\} \\
 (st)_{fv} &= (s)_{fv} \cup (t)_{fv} & (st)_{bv} &= (s)_{bv} \cup (t)_{bv} \\
 (x\langle x \rangle.t)_{fv} &= (t)_{fv} - \{x\} & (x\langle x \rangle.t)_{bv} &= (t)_{bv} \cup \{x\} \\
 (x\langle \vec{y} \rangle.t)_{fv} &= (t)_{fv} & (x\langle \vec{y} \rangle.t)_{bv} &= (t)_{bv} \\
 (u[\vec{x} \leftarrow t])_{fv} &= (u)_{fv} \cup (t)_{fv} - \{\vec{x}\} & (u[\vec{x} \leftarrow t])_{bv} &= (u)_{bv} \cup (t)_{bv} \cup \{\vec{x}\} \\
 (u[\vec{x} | y\langle y \rangle] \overline{[I]})_{fv} &= (u \overline{[I]})_{fv} - \{y\} & (u[\vec{x} | y\langle y \rangle] \overline{[I]})_{bv} &= (u \overline{[I]})_{bv} \cup \{y\}
 \end{aligned}$$

$$(u[\tilde{x} | y\langle \tilde{z} \rangle \overline{[F]}])_{fv} = (u[\overline{[F]}])_{fv} \cup \{y\} \quad (u[\tilde{x} | y\langle \tilde{z} \rangle \overline{[F]}])_{bv} = (u[\overline{[F]}])_{bv}$$

Definition 5 (Free and Bound Phantom-Variables). The **free phantom-variables** $(-)_fp$ and **bound phantom-variables** $(-)_bp$ of the pre-term t are defined as follows

$$\begin{aligned} (x)_{fp} &= \{ \} & (x)_{bp} &= \{ \} \\ (st)_{fp} &= (s)_{fp} \cup (t)_{fp} & (st)_{bp} &= (s)_{bp} \cup (t)_{bp} \\ (x\langle x \rangle.t)_{fp} &= (t)_{fp} & & \\ (c\langle \tilde{x} \rangle.t)_{fp} &= (t)_{fp} \cup \{c\} & (c\langle \tilde{x} \rangle.t)_{bp} &= (t)_{bp} \\ (u[\tilde{x} \leftarrow t])_{fp} &= (u)_{fp} \cup (t)_{fp} & (u[\tilde{x} \leftarrow t])_{bp} &= (u)_{bp} \cup (t)_{bp} \\ (u[\tilde{x} | c\langle c \rangle \overline{[F]}])_{fp} &= (u[\overline{[F]}])_{fp} - \{\tilde{x}\} & & \\ (u[\tilde{x} | c\langle \tilde{y} \rangle \overline{[F]}])_{fp} &= (u[\overline{[F]}])_{fp} \cup \{c\} - \{\tilde{x}\} & (u[\tilde{x} | c\langle \tilde{y} \rangle \overline{[F]}])_{bp} &= (u[\overline{[F]}])_{bp} \cup \{\tilde{x}\} \end{aligned}$$

The **free covers** $(u)_{fc}$ and **bound covers** $(u)_{bc}$ are the covers associated with the free phantom-variables $(u)_{fp}$ respectively the bound phantom-variables $(u)_{bp}$ of u ; that is, if x occurs as $x\langle \tilde{a} \rangle$ in u and $x \in (u)_{fp}$ then $\langle \tilde{a} \rangle \in (u)_{fc}$. When bound, x and the variables in $\tilde{a}$ may be alpha-converted independently. When a distributor $u[\tilde{x} | y\langle \tilde{z} \rangle \overline{[F]}]$ binds the phantom-variables $\tilde{x} = x_1, \dots, x_n$ where each x_i occurs as $x_i\langle \tilde{a}_i \rangle$ in u , then for technical convenience we may make the covers explicit in the distributor itself, and write

$$u[x_1\langle \tilde{a}_1 \rangle \dots x_n\langle \tilde{a}_n \rangle | y\langle \tilde{z} \rangle \overline{[F]}} .$$

The environment $\overline{[F]}$ is expected to bind *exactly* the variables in the covers $\langle \tilde{a}_i \rangle$. We apply this and other restrictions to define the terms of the calculus.

Definition 6. Terms $t \in \Lambda_a^S$ are pre-terms with the following constraints

1. Each variable may occur at most once.
2. In a phantom-abstraction $x\langle \tilde{y} \rangle.t$, $\{\tilde{y}\} \subseteq (t)_{fv}$.
3. In a sharing $u[\tilde{x} \leftarrow t]$, $\{\tilde{x}\} \subseteq (u)_{fv}$.
4. In a distributor $u[x_1\langle \tilde{a}_1 \rangle \dots x_n\langle \tilde{a}_n \rangle | y\langle \tilde{z} \rangle \overline{[F]}]$
 - (a) $\{x_1, \dots, x_n\} \subseteq (u)_{fp}$;
 - (b) the variables in $\bigcup_{i \leq n} \{\tilde{a}_i\}$ are free in u and bound by $\overline{[F]}$.
 - (c) the variables in $\{\tilde{z}\}$ occur freely in the environment $\overline{[F]}$.

Example 1. Here we show some pre-terms that are not terms.

- $c\langle x \rangle.y$ (violates condition 2)
- $xy[x, z \leftarrow w]$ (violates condition 3)
- $e_2\langle w_2 \rangle.w_2((e_1\langle w_1 \rangle.w_1)z)[e_1\langle w_1 \rangle, e_2\langle w_2 \rangle | c\langle z \rangle[w_1, w_2 \leftarrow x\langle x \rangle.xy]]$ (violates condition 4a)

We also work modulo permutation with respect to the variables in the cover of phantom-abstractions. Let $\tilde{x}$ be a list of variables and let $\tilde{x}_P$ be a permutation of that list, then the following terms are considered equal.

Fig. 5: Typing derivations for phantom-abstractions and distributors

$$u[\bar{x} \leftarrow t] \sim u[\bar{x}_P \leftarrow t] \quad y(\bar{x}).t \sim y(\bar{x}_P).t$$

Terms are typed with the typing system for Λ^S extended with the *distribution* inference rule. This rule is the result of computationally interpreting the medial rule as done in [16]. We obtain this variant of the medial rule due to the restriction for implications and to avoid introducing disjunction to the typing system. The terms of Λ_a^S are then typed as in both Figure 3 and Figure 5. Note environments are typed by the derivations of all its closures composed horizontally with the conjunction connective. Also note that in the case for phantom-abstraction is similar for that of an abstraction, where we replace one occurrence of the simple type A by the conjunction Γ .

3.1 Compilation and Readback.

We now define the translations between Λ_a^S and the original λ -calculus. First we define the interpretation $\Lambda \rightarrow \Lambda_a^S$ (*compilation*). Intuitively, it replaces each abstraction $\lambda x.-$ with the term $x\langle x \rangle.-[x_1, \dots, x_n \leftarrow x]$ where $x_1, \dots, x_n$ replace the occurrences of x . Actual substitutions are denoted as $\{t/x\}$. Let $|M|_x$ denote the number of occurrences of x in M , and if $|M|_x = n$ let $M \frac{n}{x}$ denote M with the occurrences of x replaced by fresh, distinct variables $x_1, \dots, x_n$. First, the translation of a *closed* term M is $\langle M \rangle'$, defined below

Definition 7 (Compilation). *The interpretation of λ terms, $\langle \Lambda \rangle' : \Lambda \rightarrow \Lambda_a^S$, is defined as*

$$\langle M \frac{n_1}{x_1} \dots \frac{n_k}{x_k} \rangle' [x_1^1, \dots, x_1^{n_1} \leftarrow x_1] \dots [x_k^1, \dots, x_k^{n_k} \leftarrow x_k]$$

where $x_1, \dots, x_k$ are the free variables of M such that $|M|_{x_i} = n_i > 1$ and $\langle - \rangle'$ is defined on terms as (where $n \neq 1$ in the abstraction case):

$$\begin{aligned} \langle x \rangle' &= x \\ \langle M N \rangle' &= \langle M \rangle' \langle N \rangle' \quad \langle \lambda x. M \rangle' = \begin{cases} x\langle x \rangle. \langle M \rangle' & \text{if } |M|_x = 1 \\ x\langle x \rangle. \langle M \frac{n}{x} \rangle' [x_1, \dots, x_n \leftarrow x] & \text{if } |M|_x = n \end{cases} \end{aligned}$$

The readback into the λ -calculus is slightly more complicated, specifically due to the bindings induced by the distributor. Interpreting a distributor construct as a λ -term requires (1) converting the phantom-abstractions it binds in

u into abstractions (2) collapsing the environment (3) maintaining the bindings between the converted abstractions and the intended variables located in the environment.

Definition 8. *Given a total function σ with domain D and codomain C , we **overwrite** the function with case $x \mapsto v$ where $x \in D$ and $v \in C$ such that*

$$\sigma[x \mapsto v](z) \quad := \quad \text{if } (x = z) \text{ then } v \text{ else } \sigma(z)$$

We use the map σ as part of the translation, the intuition is that for all bound variables x in the term we are translating, it should be that $\sigma(x) = x$. The purpose of the map γ is to keep track of the binding of phantom-variables.

Definition 9. *The interpretation $\llbracket - \mid - \mid - \rrbracket : \Lambda_a^S \times (V \rightarrow \Lambda) \times (V \rightarrow V) \rightarrow \Lambda$ is defined as*

$$\begin{aligned} \llbracket x \mid \sigma \mid \gamma \rrbracket &= \sigma(x) & \llbracket st \mid \sigma \mid \gamma \rrbracket &= \llbracket s \mid \sigma \mid \gamma \rrbracket \llbracket t \mid \sigma \mid \gamma \rrbracket \\ \llbracket c\langle c \rangle.t \mid \sigma \mid \gamma \rrbracket &= \lambda c. \llbracket t \mid \sigma[c \mapsto c] \mid \gamma \rrbracket \\ \llbracket c\langle x_1, \dots, x_n \rangle.t \mid \sigma \mid \gamma \rrbracket &= \lambda c. \llbracket t \mid \sigma[x_i \mapsto \sigma(x_i)\{c/\gamma(c)\}]_{i \in [n]} \mid \gamma \rrbracket \\ \llbracket u[x_1, \dots, x_n \leftarrow t] \mid \sigma \mid \gamma \rrbracket &= \llbracket u \mid \sigma[x_i \mapsto \llbracket t \mid \sigma \mid \gamma \rrbracket]_{i \in [n]} \mid \gamma \rrbracket \\ \llbracket u[e_1\langle \bar{w}_1 \rangle, \dots, e_n\langle \bar{w}_n \rangle \mid c\langle c \rangle[\overline{F}]] \mid \sigma \mid \gamma \rrbracket &= \llbracket u[\overline{F}] \mid \sigma \mid \gamma[e_i \mapsto c]_{i \in [n]} \rrbracket \\ \llbracket u[e_1\langle \bar{w}_1 \rangle, \dots, e_n\langle \bar{w}_n \rangle \mid c\langle x_1, \dots, x_m \rangle[\overline{F}]] \mid \sigma \mid \gamma \rrbracket &= \llbracket u[\overline{F}] \mid \sigma' \mid \gamma[e_i \mapsto c]_{i \in [n]} \rrbracket \end{aligned}$$

where $\sigma' = \sigma[x_i \mapsto \sigma(x_i)\{c/\gamma(c)\}]_{i \in [n]}$

The following Proposition justifies working modulo permutation equivalence.

Proposition 1. *For $s, t \in \Lambda_a^S$, if $s \sim t$ then $\llbracket s \rrbracket = \llbracket t \rrbracket$.*

3.2 Rewrite Rules.

Both the spinal atomic λ -calculus and the atomic λ -calculus of [16] follow atomic reduction steps, i.e. they apply on individual constructors. The biggest difference is that our calculus is capable of duplicating not only the skeleton but also the spine. The rewrite rules in our calculus make use of 3 operations, *substitution*, *book-keeping*, and *exorcism*. The operation **substitution** $t\{s/x\}$ propagates through the term t , and replaces the free occurrences of the variable x with the term s . Moreover, if x occurs in the cover of a phantom-variable $e\langle \bar{y} \cdot x \rangle$, then substitution replaces the x in the cover with $(s)_{fv}$, resulting in $e\langle \bar{y} \cdot (s)_{fv} \rangle$. Although substitution performs some book-keeping on phantom-abstractions, we define an explicit notion of **book-keeping** $\{\bar{y}/e\}_b$ that updates the variables stored in a free cover i.e. for a term t , $e\langle \bar{x} \rangle \in (t)_{fc}$ then $e\langle \bar{y} \rangle \in (t\{\bar{y}/e\}_b)_{fc}$. The last operation we introduce is called **exorcism** $\{c\langle \bar{x} \rangle\}_e$. We perform exorcisms on phantom-abstractions to convert them to abstractions. Intuitively, this will be performed on phantom-abstractions with phantom-variables bound to a distributor when said distributor is eliminated. It converts phantom-abstractions to abstractions by introducing a sharing of the phantom-variable that captures the variables in the cover, i.e. $(c\langle \bar{x} \rangle.t)\{c\langle \bar{x} \rangle\}_e = c\langle c \rangle.t[\bar{x} \leftarrow c]$.

Proposition 2. *The translation $\llbracket u \mid \sigma \mid \gamma \rrbracket$ commutes with substitutions, book-keepings¹, and exorcisms² in the following way*

$$\llbracket u\{t/x\} \mid \sigma \mid \gamma \rrbracket = \llbracket u \mid \sigma[x \mapsto \llbracket t \mid \sigma \mid \gamma \rrbracket] \mid \gamma \rrbracket$$

$$\llbracket u\{\bar{x}/c\}_b \mid \sigma \mid \gamma \rrbracket = \llbracket u \mid \sigma \mid \gamma \rrbracket$$

$$\llbracket u\{c\langle x_1, \dots, x_n \rangle\}_e \mid \sigma \mid \gamma \rrbracket = \llbracket u \mid \sigma[x_i \mapsto c]_{i \in [n]} \mid \gamma \rrbracket$$

(1) Given $c\langle \bar{y} \rangle \in (u)_{fc}$ where $\bar{x} \subseteq \bar{y}$ and for $z \in \bar{y}/\bar{x}$, $\gamma(c) \notin (\sigma(z))_{fv}$

(2) Given $c\langle \bar{x} \rangle \in (u)_{fc}$ or $\{\bar{x}\} \cap (u)_{fv} = \{\}$

Proof. See [25], proof of Proposition 18, 19, 20, 21.

Using these operations, we define the rewrite rules that allow for spinal duplication. Firstly we have beta reduction ($\rightsquigarrow_\beta$), which strictly requires an abstraction (not a phantom).

$$(x\langle x \rangle.t)s \rightsquigarrow_\beta t\{s/x\}$$

$$\frac{\Gamma}{A \rightarrow \boxed{A^x \wedge \Gamma \parallel_t B}} \lambda \quad \wedge \quad \boxed{\Delta \parallel_s A}$$

$$\rightsquigarrow_\beta$$

$$\frac{\boxed{\Delta \parallel_s A} \wedge \Gamma}{\boxed{A \wedge \Gamma \parallel_t B}} \textcircled{\wedge}$$

(β)

Here β -reduction is a linear operation, since the bound variable x occurs exactly once in the body t . Any duplication of the term t in the atomic λ -calculus proceeds via the sharing reductions.

The first set of sharing reduction rules move closures towards the outside of a term. Most of these rewrite rules only change the typing derivations in the way that subderivations are composed, with the exception of moving a closure out of scope of a distributor.

$$s[\Gamma]t \rightsquigarrow_L (st)[\Gamma] \quad (l_1)$$

$$st[\Gamma] \rightsquigarrow_L (st)[\Gamma] \quad (l_2)$$

$$d\langle \bar{x} \rangle.t[\Gamma] \rightsquigarrow_L (d\langle \bar{x} \rangle.t)[\Gamma] \text{ if } \{\bar{x}\} \cap (t)_{fv} = \{\bar{x}\} \quad (l_3)$$

$$u[\bar{x} \leftarrow t[\Gamma]] \rightsquigarrow_L u[\bar{x} \leftarrow t][\Gamma] \quad (l_4)$$

For the case of lifting a closure outside a distributor, we use a notation $\parallel [\Gamma] \parallel$ to identify the variables captured by a closure, i.e. $\parallel [\bar{x} \leftarrow t] \parallel = \{\bar{x}\}$ and $\parallel [e_1\langle \bar{x}_1 \rangle, \dots, e_n\langle \bar{x}_n \rangle \mid c\langle c \rangle[\Gamma]] \parallel = \{\bar{x}_1, \dots, \bar{x}_n\}$. Then let $\{\bar{z}\} = \parallel [\Gamma] \parallel$ in the following rewrite rule, where we remove $\bar{z}$ from the covers, that can only occur if $\{\bar{x}\} \cap ([\Gamma])_{fv} = \{\}$.

$$\begin{aligned} & u[e_1\langle \bar{w}_1 \rangle \dots e_n\langle \bar{w}_n \rangle \mid c\langle \bar{x} \rangle[\Gamma][\Gamma]] \\ & \rightsquigarrow_L u\{(\bar{w}_i \setminus \bar{z})/e_i\}_{b_i \in [n]}[e_1\langle \bar{w}_1 \setminus \bar{z} \rangle \dots e_n\langle \bar{w}_n \setminus \bar{z} \rangle \mid c\langle \bar{x} \rangle[\Gamma]]][\Gamma] \end{aligned} \quad (l_5)$$

The graphical version of this rule is shown in Figure 4c, where we remove the edge only if there is no edge between t and the unsharing node. The proof rewrite rule corresponding with the rewrite rule l_5 can be broken down into two parts. The first part is readjusting how the derivations compose as shown below.

$$\begin{array}{c}
 \frac{(C \rightarrow \Gamma) \wedge \Delta \wedge \Omega}{C \rightarrow} \quad \frac{\Gamma \wedge \Delta \wedge \frac{\Omega}{A \wedge \dots \wedge A}}{\Sigma_1 \dots \Sigma_n} \quad \frac{(C \rightarrow \Sigma_1) \wedge \dots \wedge (C \rightarrow \Sigma_n)}{(C \rightarrow \Gamma) \wedge \Delta \wedge \Omega} \quad s \\
 \sim_L \quad \frac{(C \rightarrow \Gamma) \wedge \Delta \wedge \frac{\Omega}{A \wedge \dots \wedge A}}{C \rightarrow} \quad \frac{\Gamma \wedge \Delta \wedge A \dots A}{\Sigma_1 \dots \Sigma_n} \quad \frac{(C \rightarrow \Sigma_1) \wedge \dots \wedge (C \rightarrow \Sigma_n)}{(C \rightarrow \Gamma) \wedge \Delta \wedge A \dots A} \quad d
 \end{array}$$

The second part of the rewrite rule justifies the need for the book-keeping operation. In the rewrite below, let A be the type of a variable z where $z \in \bar{z}$. After lifting, we want to remove the variable from the cover as to ensure correctness since the variables in the cover denote the variables captured by the environment. Book-keeping allows us to remove these variables simultaneously.

$$\begin{array}{c}
 \frac{(C \rightarrow \Gamma) \wedge \Delta \wedge A}{C \rightarrow} \quad \frac{\Gamma \wedge \Delta}{\Sigma_1 \wedge \dots \wedge \Sigma_n} \quad \frac{(C \rightarrow \Sigma_1) \wedge \dots \wedge (C \rightarrow \Sigma_n)}{\dots \wedge (C \rightarrow \Sigma_i \wedge A) \wedge \dots} \quad d \\
 \sim \quad \frac{(C \rightarrow \Gamma) \wedge \Delta}{C \rightarrow} \quad \frac{\Gamma \wedge \Delta}{\Sigma_1 \wedge \dots \wedge \Sigma_n} \quad \frac{(C \rightarrow \Sigma_1) \wedge \dots \wedge (C \rightarrow \Sigma_i) \wedge \dots}{\dots \wedge (C \rightarrow \Sigma_i \wedge A) \wedge \dots} \quad d \quad \wedge A
 \end{array}$$

The lifting rules (l_i) are justified by the need to lift closures out of the distributor, as opposed to duplicating them. The second set of rewrite rules, consecutive sharings are compounded and unary sharings are applied as substitutions. For simplicity, in the equivalent proof rewrite step we only show the binary case.

$$u[\bar{w} \leftarrow y][y \cdot \bar{y} \leftarrow t] \rightsquigarrow_C u[\bar{w} \cdot \bar{y} \leftarrow t] \quad (c_1)$$

$$u[x \leftarrow t] \rightsquigarrow_C u\{t/x\} \quad (c_2)$$

$$\frac{A}{A \wedge \frac{A}{A \wedge A} \Delta} \Delta \quad \rightsquigarrow_C \quad \frac{A}{A \wedge A \wedge A} \Delta \quad \frac{A}{A} \Delta \quad \rightsquigarrow_C \quad A$$

The atomic steps for duplicating are given in the third and final set of rewrite rules. The first being the atomic duplication step of an application, which is the same rule used in [16]. The binary case proof rewrite steps for each rule are also provided. There are also shown graphically in (respectively) Figure 4b (where we maintain links between sharings and unsharings), Figure 4a, and Figure 4d (where the unsharing node is linked to exactly one connecting sharing node).

$$u[x_1 \dots x_n \leftarrow st] \rightsquigarrow_D u\{z_1 y_1/x_1\} \dots \{z_n y_n/x_n\}[z_1 \dots z_n \leftarrow s][y_1 \dots y_n \leftarrow t] \quad (d_1)$$

$$\begin{array}{c}
\frac{(A \rightarrow B) \wedge A}{\frac{B}{B \wedge B} \Delta} @ \quad \rightsquigarrow_D \quad \frac{\frac{(A \rightarrow B)}{(A \rightarrow B) \wedge (A \rightarrow B)} \Delta \wedge \frac{B}{B \wedge B} \Delta}{\frac{(A \rightarrow B) \wedge A}{B} @ \wedge \frac{(A \rightarrow B) \wedge A}{B} @} \\
\\
u[x_1, \dots, x_n \leftarrow c(\bar{y}).t] \rightsquigarrow_D \\
u\{e_i\langle w_i \rangle.w_i/x_i\}_{i \in [n]}[e_1\langle w_1 \rangle \dots e_n\langle w_n \rangle \mid c(\bar{y})[w_1, \dots, w_n \leftarrow t]] \quad (d_2)
\\
\\
\frac{\frac{(A \rightarrow B) \wedge \Gamma}{A \rightarrow \boxed{\frac{B \wedge \Gamma}{\parallel} C}} s}{(A \rightarrow C) \wedge (A \rightarrow C) \Delta} \rightsquigarrow_D \frac{\frac{(A \rightarrow B) \wedge \Gamma}{A \rightarrow \boxed{\frac{B \wedge \Gamma}{\parallel} C}} s}{\frac{\frac{B \wedge \Gamma}{\parallel} C}{C \wedge C} \Delta} d \\
\\
u[e_1\langle \bar{w}_1 \rangle \dots e_n\langle \bar{w}_n \rangle \mid c\langle c \rangle[\bar{w}_1, \dots, \bar{w}_n \leftarrow c]] \rightsquigarrow_D u\{e_1\langle \bar{w}_1 \rangle\}_e \dots \{e_n\langle \bar{w}_n \rangle\}_e \quad (d_3)
\\
\\
\frac{\frac{\frac{A}{A \wedge A} \Delta}{(A \rightarrow A) \wedge (A \rightarrow A)} a}{\frac{A \rightarrow A}{A \wedge A} \Delta} \rightsquigarrow_D \frac{A \rightarrow A}{A \wedge A} a \wedge \frac{A \rightarrow A}{A \wedge A} a
\end{array}$$

Example 2. The following example, illustrated in Figure 4e, is a reduction in the term calculus where we duplicate the spine of the term $[a_1, a_2 \leftarrow \underline{\lambda x. \lambda y. ((\lambda z. z)y) \underline{x}}]$.

$$\begin{array}{l}
\rightsquigarrow_D \{ \underline{x_1}\langle b_1 \rangle. b_1/a_1 \} \{ \underline{x_2}\langle b_2 \rangle. b_2/a_2 \} [x_1\langle b_1 \rangle, x_2\langle b_2 \rangle \mid x\langle x \rangle[b_1, b_2 \leftarrow \underline{\lambda y. ((\lambda z. z)y) \underline{x}}]] \\
\rightsquigarrow_D \{ \underline{x_1}\langle c_1 \rangle. \underline{y_1}\langle c_1 \rangle c_1/a_1 \} \{ \underline{x_2}\langle c_2 \rangle. \underline{y_2}\langle c_2 \rangle. c_2/a_2 \} \\
\quad [x_1\langle c_1 \rangle, x_2\langle c_2 \rangle \mid x\langle x \rangle[y_1\langle c_1 \rangle, y_2\langle c_2 \rangle \mid y\langle y \rangle[c_1, c_2 \leftarrow \underline{((\lambda z. z)y) \underline{x}}]]] \\
\rightsquigarrow_D \{ \underline{x_1}\langle d_1, e_1 \rangle. \underline{y_1}\langle d_1, e_1 \rangle d_1 e_1/a_1 \} \{ \underline{x_2}\langle d_2, e_2 \rangle. \underline{y_2}\langle d_2, e_2 \rangle. d_2 e_2/a_2 \} \\
\quad [x_1\langle d_1, e_1 \rangle, x_2\langle d_2, e_2 \rangle \mid x\langle x \rangle[y_1\langle d_1, e_1 \rangle, y_2\langle d_2, e_2 \rangle \mid y\langle y \rangle[d_1, d_2 \leftarrow (\lambda z. z)y][e_1, e_2 \leftarrow \underline{x}]]] \\
\rightsquigarrow_L \{ \underline{x_1}\langle d_1, e_1 \rangle. \underline{y_1}\langle d_1 \rangle d_1 e_1/a_1 \} \{ \underline{x_2}\langle d_2, e_2 \rangle. \underline{y_2}\langle d_2 \rangle. d_2 e_2/a_2 \} \\
\quad [x_1\langle d_1, e_1 \rangle, x_2\langle d_2, e_2 \rangle \mid x\langle x \rangle[y_1\langle d_1 \rangle, y_2\langle d_2 \rangle \mid y\langle y \rangle[d_1, d_2 \leftarrow (\lambda z. z)y][e_1, e_2 \leftarrow \underline{x}]]] \\
\rightsquigarrow_L \{ \underline{x_1}\langle e_1 \rangle. \underline{y_1}\langle d_1 \rangle d_1 e_1/a_1 \} \{ \underline{x_2}\langle e_2 \rangle. \underline{y_2}\langle d_2 \rangle. d_2 e_2/a_2 \} \\
\quad [x_1\langle e_1 \rangle, x_2\langle e_2 \rangle \mid x\langle x \rangle[e_1, e_2 \leftarrow \underline{x}]] [y_1\langle d_1 \rangle, y_2\langle d_2 \rangle \mid y\langle y \rangle[d_1, d_2 \leftarrow (\lambda z. z)y]] \\
\rightsquigarrow_D \{ \underline{\lambda x_1. \underline{y_1}}\langle d_1 \rangle d_1 \underline{x_1}/a_1 \} \{ \underline{\lambda x_2. \underline{y_2}}\langle d_2 \rangle. d_2 \underline{x_2}/a_2 \} [y_1\langle d_1 \rangle, y_2\langle d_2 \rangle \mid y\langle y \rangle[d_1, d_2 \leftarrow (\lambda z. z)y]]
\end{array}$$

Reduction ($\rightsquigarrow_{(L,C,D,\beta)}$) preserves the conclusion of the derivation, and thus the following proposition is easy to observe.

Proposition 3. *If $s \rightsquigarrow_{(L,C,D,\beta)} t$ and $s : A$, then $t : A$.*

Definition 10. *For a term $t \in \Lambda_a^S$, if there does not exists a term $s \in \Lambda_a^S$ such that $t \rightsquigarrow_{(L,C,D)} s$ then it is said that t is in **sharing normal form**.*

The following Lemma not only proves we have good translations in Section 3.1, and shows duplication preserves denotation.

Lemma 1. For a $t \in \Lambda_a^S$ in sharing normal form and a $N \in \Lambda$.

$$\llbracket \langle N \rangle \rrbracket = N \quad \llbracket \langle t \rangle \rrbracket = t \quad \exists_{M \in \Lambda}. t = \langle M \rangle$$

Otherwise if $s \rightsquigarrow_{(L,D,C)} t$ then $\llbracket s \mid \sigma \mid \gamma \rrbracket = \llbracket t \mid \sigma \mid \gamma \rrbracket$.

Proof. See [25, Lemma 24, Lemma 25].

Lemma 2. Given a term $t \in \Lambda_a^S$, then $\llbracket \langle t \rangle \rrbracket$ is t in sharing normal form.

Proof. We can prove this by induction on the longest sharing reduction path from t . Our base case is already covered by Lemma 1. We are then interested in the inductive case, where t is not in sharing normal form. By Lemma 1, $\llbracket t \rrbracket = \llbracket t' \rrbracket$ where $t \rightsquigarrow_{(D,L,C)} t'$. By induction hypothesis, $\llbracket \langle t' \rangle \rrbracket$ is in sharing normal form. Hence $\llbracket \langle t \rangle \rrbracket$ is in sharing normal form. $\square$

4 Strong Normalisation of Sharing Reductions

In order to show our calculus is strongly normalising, we first show that the sharing reduction rules are strongly normalising. We induct a measure on terms and show that this measure strictly decreases as sharing reduction progresses. Similar ideas and results can be found elsewhere: with *memory* in [20], the λ -*I* calculus in [6], the λ -void calculus [2], and the weakening $\lambda\mu$ -calculus [17]. Our measure will consist of three components. First, the **height** of a term is a multiset of integers, that measures the number of constructors from each sharing node to the root of the term in its graphical notation. The height is defined on terms as $\mathcal{H}^i(-)$, where i is an integer. We say $\mathcal{H}(t)$ for $\mathcal{H}^1(t)$. We use $\cup$ to denote the disjoint union of two multisets. We denote $\mathcal{H}^i([\Gamma_1]) \cup \dots \cup \mathcal{H}^i([\Gamma_n])$ as $\mathcal{H}^i([\overline{\Gamma}])$ for the environment $[\overline{\Gamma}] = [\Gamma_1], \dots, [\Gamma_n]$.

Definition 11 (Sharing Height). The sharing height $\mathcal{H}^i(t)$ of a term t is given below, where n is the number of closures in $[\overline{\Gamma}]$:

$$\begin{aligned} \mathcal{H}^i(x) &= \{\} & \mathcal{H}^i(st) &= \mathcal{H}^{i+1}(s) \cup \mathcal{H}^{i+1}(t) \\ \mathcal{H}^i(c\langle \vec{x} \rangle.t) &= \mathcal{H}^{i+1}(t) & \mathcal{H}^i(t[\Gamma]) &= \mathcal{H}^i(t) \cup \mathcal{H}^i([\Gamma]) \cup \{i^1\} \\ \mathcal{H}^i([x_1, \dots, x_n \leftarrow t]) &= \mathcal{H}^{i+1}(t) & \mathcal{H}^i([\vec{w} \mid c\langle \vec{x} \rangle [\overline{\Gamma}]] &= \mathcal{H}^{i+1}([\overline{\Gamma}]) \cup \{(i+1)^n\} \end{aligned}$$

This measure then strictly decreases for the rewrite rules l_1 , l_2 , l_3 , l_4 and l_5 , i.e. if $t \rightsquigarrow_L u$ then $\mathcal{H}^i(t) > \mathcal{H}^i(u)$. The second measure we consider is the **weight** of a term. Intuitively this quantifies the remaining duplications, which are performed with $\rightsquigarrow_D$ reductions. If a term would be deleted, we assign it with a weight ‘1’ to express that it is not duplicated. Calculating the weight requires an auxiliary function that assigns integer weights to the variables of a term. This function is defined on terms $\mathcal{V}^i(-)$, where i is an integer. To measure variables independently of binders is vital. It allows to measure distributors, which duplicate λ ’s but not the bound variable. Also, only bound variables for abstractions are measured since variables bound by sharings are substituted in the interpretation.

Definition 12 (Variable Weights). *The function $\mathcal{V}^i(t)$ returns a function that assigns integer weights to the free variables of t . It is defined by the below, where $f = \mathcal{V}^i(t)$ and $g = f(x_1) + \dots + f(x_n)$ for each $x_i \in \vec{x}$.*

$$\begin{aligned}
\mathcal{V}^i(x) &= \{x \mapsto i\} & \mathcal{V}^i(st) &= \mathcal{V}^i(s) \cup \mathcal{V}^i(t) \\
\mathcal{V}^i(c\langle c \rangle.t) &= \mathcal{V}^i(t)/\{c\} & \mathcal{V}^i(c\langle \vec{x} \rangle.t) &= \mathcal{V}^i(t) \cup \{c \mapsto i\} \\
\mathcal{V}^i(t[\vec{x} \leftarrow s]) &= \mathcal{V}^i(t)/\{\vec{x}\} \cup \mathcal{V}^g(s) & \mathcal{V}^i(t[\leftarrow s]) &= \mathcal{V}^i(t) \cup \mathcal{V}^1(s) \\
\mathcal{V}^i(t[e_1\langle \vec{w}_1 \rangle \dots e_n\langle \vec{w}_n \rangle | c\langle c \rangle \overline{[F]}]) &= \mathcal{V}^i(t[\overline{[F]}])/\{c, e_1, \dots, e_n\} \\
\mathcal{V}^i(t[e_1\langle \vec{w}_1 \rangle \dots e_n\langle \vec{w}_n \rangle | c\langle \vec{x} \rangle \overline{[F]}]) &= \mathcal{V}^i(t[\overline{[F]}])/\{e_1, \dots, e_n\} \cup \{c \mapsto i\}
\end{aligned}$$

The weight of a term can then be defined via the use of this auxiliary function. The auxiliary function is used when calculating the weight of a sharing, where the sharing weight of the variables bound by the sharing play a significant role in calculating the weight of the shared term. In the case of a weakening $[\leftarrow t]$, we assign an initial weight of 1. Again we say $\mathcal{W}(t) = \mathcal{W}^1(t)$.

Definition 13 (Sharing Weight). *The sharing weight $\mathcal{W}^i(t)$ of a term t is a multiset of integers computed by the function defined below, where $f = \mathcal{V}^i(t)$ and $g = f(x_1) + \dots + f(x_n)$ for each $x_i \in \vec{x}$.*

$$\begin{aligned}
\mathcal{W}^i(x) &= \{\} & \mathcal{W}^i(st) &= \mathcal{W}^i(s) \cup \mathcal{W}^i(t) \cup \{i\} \\
\mathcal{W}^i(c\langle c \rangle.t) &= \mathcal{W}^i(t) \cup \{i\} \cup \{\mathcal{V}^i(t)(c)\} & \mathcal{W}^i(c\langle \vec{x} \rangle.t) &= \mathcal{W}^i(t) \cup \{i\} \\
\mathcal{W}^i(t[\vec{x} \leftarrow s]) &= \mathcal{W}^i(t) \cup \mathcal{W}^g(s) & \mathcal{W}^i(t[\leftarrow s]) &= \mathcal{W}^i(t) \cup \mathcal{W}^1(s) \\
\mathcal{W}^i(t[e_1\langle \vec{w}_1 \rangle \dots e_n\langle \vec{w}_n \rangle | c\langle c \rangle \overline{[F]}]) &= \mathcal{W}^i(t[\overline{[F]}]) \cup \{\mathcal{V}^i(t[\overline{[F]}])(c)\} \\
\mathcal{W}^i(t[e_1\langle \vec{w}_1 \rangle \dots e_n\langle \vec{w}_n \rangle | c\langle \vec{x} \rangle \overline{[F]}]) &= \mathcal{W}^i(t[\overline{[F]}])
\end{aligned}$$

This measure then strictly decreases on the rewrite rules d_1, d_2, d_3 and is unaffected by all the other sharing reduction rules, i.e. if $t \rightsquigarrow_D u$ then $\mathcal{W}^i(t) > \mathcal{W}^i(u)$. If $t \rightsquigarrow_{(L,C)} u$ then $\mathcal{W}^i(t) = \mathcal{W}^i(u)$. The third and last measure we consider is the **number of closures** in the term, where it can be easily observed that the rewrite rules c_1 and c_2 strictly decrease this measure, and that the $\rightsquigarrow_L$ rules do not alter the number of closures. We then use this along with height and weight to define a *sharing measure* on terms.

Definition 14. *The **sharing measure** of a Λ_a^S -term t is a triple $(\mathcal{W}(t), \mathcal{C}, \mathcal{H}(t))$, where $\mathcal{C}$ is the number of closures in the term t . We compare sharing measures by using the lexicographical preferences according to $\mathcal{W} > \mathcal{C} > \mathcal{H}$.*

Theorem 1. *Sharing reduction $\rightsquigarrow_{(D,L,C)}$ is strongly normalising.*

Now that we have proven the sharing reductions are strongly normalising, we can prove that they are confluent for closed terms.

Theorem 2. *The sharing reduction relation $\rightsquigarrow_{(D,L,C)}$ is confluent.*

Proof. Lemma 1 tells us that the preservation is preserved under reduction i.e. for $s \rightsquigarrow_{(D,L,C)} t$, $\llbracket s \rrbracket = \llbracket t \rrbracket$. Therefore given $t \rightsquigarrow_{(D,L,C)}^* s_1$ and $t \rightsquigarrow_{(D,L,C)}^* s_2$, $\llbracket t \rrbracket = \llbracket s_1 \rrbracket = \llbracket s_2 \rrbracket$. Since we know that sharing reductions are strongly normalising, we know there exists terms u_1 and u_2 in sharing normal form such that $s_1 \rightsquigarrow_{(D,L,C)}^* u_1$ and $s_2 \rightsquigarrow_{(D,L,C)}^* u_2$. Lemma 1 tells us that terms in sharing normal form are in correspondence with their denotations i.e. $\llbracket \llbracket t \rrbracket \rrbracket = t$. Since by Lemma 1 we know $\llbracket u_1 \rrbracket = \llbracket s_1 \rrbracket = \llbracket s_2 \rrbracket = \llbracket u_2 \rrbracket$, and by Lemma 1 $\llbracket \llbracket u_1 \rrbracket \rrbracket = u_1$ and $\llbracket \llbracket u_2 \rrbracket \rrbracket = u_2$, we can conclude $u_1 = u_2$. Hence, we prove confluence. $\square$

5 Preservation of Strong Normalisation and Confluence

A β -step in our calculus may occur within a weakening, and therefore is simulated by zero β -steps in the λ -calculus. Therefore if there is an infinite reduction path located inside a weakening in Λ_a^S , then the reduction path is not preserved in the corresponding λ -term as there are no weakenings. To deal with this, just as done in [2, 16, 17], we make use of the **weakening calculus**. A β -step is non-deleting precisely because of the weakening construct. If a β -step would be deleting, then the weakening calculus would instead keep the deleted term around as ‘garbage’, which can continue to reduce unless explicitly ‘garbage-collected’ by extra (non- β) reduction steps. PSN has already been shown for the weakening calculus through the use of a perpetual strategy in [16]. A part of proving PSN is then using the weakening calculus to prove that if $t \in \Lambda_a^S$ has a infinite reduction path, then its translation into the weakening calculus also has an infinite reduction path.

Definition 15. *The w -terms of the weakening calculus (Λ_w) are*

$$T, U, V ::= x \mid \lambda x. T^* \mid UV \mid T[\leftarrow U] \mid \bullet (*) \text{ where } x \in (T)_{fv}$$

The terms are variable, abstraction, application, weakening, and a bullet. In the weakening $T[\leftarrow U]$, the subterm U is *weakened*. The interpretation of atomic terms to weakening terms $\llbracket - \mid - \mid - \rrbracket_w$ can be seen as an extension of the translation into the λ -calculus (Definition 9).

Definition 16. *The interpretation $\llbracket - \mid - \mid - \rrbracket_w : \Lambda_a^S \times (V \rightarrow \Lambda_w) \times (V \rightarrow V) \rightarrow \Lambda_w$ with maps $\sigma : V \rightarrow \Lambda_w$ and $\gamma : V \rightarrow V$ is defined as an extension of the translation in (Definition 9) with the following additional special cases.*

$$\begin{aligned} \llbracket u[\leftarrow t] \mid \sigma \mid \gamma \rrbracket_w &= \llbracket u \mid \sigma \mid \gamma \rrbracket_w [\leftarrow \llbracket t \mid \sigma \mid \gamma \rrbracket_w] \\ \llbracket u[\mid c(c) \mid \overline{[T]}] \mid \sigma \mid \gamma \rrbracket_w &= \llbracket u[\overline{[T]}] \mid \sigma[c \mapsto \bullet] \mid \gamma \rrbracket_w \\ \llbracket u[\mid c(x_1, \dots, x_n) \mid \overline{[T]}] \mid \sigma \mid \gamma \rrbracket_w &= \llbracket u[\overline{[T]}] \mid \sigma' \mid \gamma \rrbracket_w \end{aligned}$$

where $\sigma'(z) :=$ if $z \in \{x_1, \dots, x_n\}$ then $\sigma(z)\{\bullet/\gamma(c)\}$ else $\sigma(z)$

We say $\llbracket t \rrbracket^w = \llbracket t \mid I \mid I \rrbracket_w$ where I is the identity function. We also have translations of the weakening calculus to and from the λ -calculus. Both of these translations were provided in [16]. The interpretation $\llbracket - \rrbracket$ from weakening terms to λ -terms discards all weakenings.

Definition 17. The interpretation $M \in \Lambda$, $\llbracket - \rrbracket^w : \Lambda \rightarrow \Lambda_w$ is defined below.

$$\llbracket x \rrbracket^w = x \quad \llbracket M N \rrbracket^w = \llbracket M \rrbracket^w \llbracket N \rrbracket^w \quad \llbracket \lambda x. N \rrbracket^w = \begin{cases} \lambda x. \llbracket N \rrbracket^w & \text{if } x \in (N)_{fv} \\ \lambda x. \llbracket N \rrbracket^w [\leftarrow x] & \text{otherwise} \end{cases}$$

The following equalities can be observed, where $\sigma^A(z) = \lfloor \sigma^w(z) \rfloor$.

Proposition 4. For $N \in \Lambda$ and $t \in \Lambda_a^S$ the following properties hold

$$\lfloor \llbracket t \mid \sigma^w \mid \gamma \rrbracket^w \rfloor = \llbracket t \mid \sigma^A \mid \gamma \rrbracket \quad \llbracket \llbracket N \rrbracket^w \rrbracket^w = \llbracket N \rrbracket^w \quad \lfloor \llbracket N \rrbracket^w \rfloor = N$$

where for each $\{x \mapsto M\} \in \sigma^w$, $\{x \mapsto \lfloor M \rfloor\} \in \sigma^A$.

Definition 18. In the weakening calculus, β -reduction is defined as follows, where $\overline{[T]}$ are weakening constructs. $((\lambda x. T) \overline{[T]}) U \rightarrow_\beta T \{U/x\} \overline{[T]}$

Proposition 5. If $N \in \Lambda$ is strongly normalising, then so is $\llbracket N \rrbracket^w$.

When translating from Λ_a^S to Λ_w , weakenings are maintained whilst sharings are interpreted via substitution. Thus the reduction rules in the weakening calculus cover the spinal reductions for nullary distributors and weakenings.

Definition 19. Weakening reduction ($\rightarrow_w$) proceeds as follows.

$$\begin{array}{ll} U[\leftarrow T]V \rightarrow_w (UV)[\leftarrow T] & UV[\leftarrow T] \rightarrow_w (UV)[\leftarrow T] \\ T[\leftarrow U[\leftarrow V]] \rightarrow_w T[\leftarrow U][\leftarrow V] & T[\leftarrow \lambda x. U] \rightarrow_w T[\leftarrow U\{\bullet/x\}] \\ T[\leftarrow UV] \rightarrow_w T[\leftarrow U][\leftarrow V] & T[\leftarrow \bullet] \rightarrow_w T \\ T[\leftarrow U] \rightarrow_w T^{(1)} & \lambda x. T[\leftarrow U] \rightarrow_w (\lambda x. T)[\leftarrow U]^{(2)} \end{array}$$

(1) if U is a subterm of T and (2) if $x \notin (U)_{fv}$

It is easy to see that these rules correspond to special cases of the sharing reduction rules for Λ_a^S . This resemblance is confirmed by the following Lemma, proven in [25, pp. 82-86]. We use this to show how Λ_a^S enjoys PSN.

Lemma 3. If $t \rightsquigarrow_\beta u$ then $\llbracket t \rrbracket^w \rightarrow_\beta^+ \llbracket u \rrbracket^w$. If $t \rightsquigarrow_{(C,D,L)} u$ and for any $x \in (t)_{bv} \cup (t)_{fp}$ such that for all z , $x \notin (\sigma(z))_{fv}$.

$$\llbracket t \mid \sigma \mid \gamma \rrbracket^w \rightarrow_w^* \llbracket u \mid \sigma \mid \gamma \rrbracket^w$$

Lemma 4. For $t \in \Lambda_a^S$ has an infinite reduction path, then $\llbracket t \rrbracket^w$ also has an infinite reduction path.

Proof. Due to Theorem 2, we know that the infinite reduction path contains infinite β -steps. This means in the reduction sequence, between each β -step, there are finite many $\rightsquigarrow_{(D,L,C)}$ reduction steps. Lemma 3 says each $\rightsquigarrow_{(D,L,C)}$ step in Λ_a^S corresponds to zero or more weakening reductions ($\rightsquigarrow_w^*$). Lemma 3 says that each beta step in Λ_a^S corresponds to one or more β -steps in Λ_w . Therefore, it must be that $\llbracket t \rrbracket^w$ also has an infinite reduction path. $\square$

Theorem 3. *If $N \in \Lambda$ is strongly normalising, then so is $\llbracket N \rrbracket$.*

Proof. For a given $N \in \Lambda$ that is strongly normalising, we know by Lemma 5 that $\llbracket N \rrbracket^\omega$ is strongly normalising. Then $\llbracket \llbracket N \rrbracket \rrbracket^\omega$ is strongly normalising, since Proposition 4 states that $\llbracket N \rrbracket^\omega = \llbracket \llbracket N \rrbracket \rrbracket^\omega$. Then by Lemma 4, which states that if $\llbracket t \rrbracket^\omega$ is strongly normalising, then t is strongly normalising, proves that $\llbracket N \rrbracket$ is strongly normalising. $\square$

We also prove confluence, which is already known for the λ -calculus [11]. We first observe that a β -step in the λ -calculus is simulated in Λ_a^S by one β -step followed by zero or more sharing reductions.

Lemma 5. *Given $N, M \in \Lambda$. If $N \rightarrow_\beta M$, then $\llbracket N \rrbracket \rightarrow_{(\beta, D, L, C)}^* \llbracket M \rrbracket$.*

Proof. This is proven by Sherratt in [25, Lemma 67].

Theorem 4. *Given $t, s_1, s_2 \in \Lambda_a^S$. If $t \rightarrow_{(\beta, D, L, C)}^* s_1$ and $t \rightarrow_{(\beta, D, L, C)}^* s_2$, there exists a $u \in \Lambda_a^S$ such that $s_1 \rightarrow_{(\beta, D, L, C)}^* u$ and $s_2 \rightarrow_{(\beta, D, L, C)}^* u$.*

Proof. Suppose $t \rightarrow_{(\beta, D, L, C)}^* s_1$ and $t \rightarrow_{(\beta, D, L, C)}^* s_2$. Then we have $\llbracket t \rrbracket \rightarrow_\beta^* \llbracket s_1 \rrbracket$ and $\llbracket t \rrbracket \rightarrow_\beta^* \llbracket s_2 \rrbracket$. By the Church-Rosser theorem, there exists a $M \in \Lambda$ such that $\llbracket s_1 \rrbracket \rightarrow_\beta^* M$ and $\llbracket s_2 \rrbracket \rightarrow_\beta^* M$. Due to Lemma 2, $\llbracket \llbracket s_1 \rrbracket \rrbracket = s'_1$ and $\llbracket \llbracket s_2 \rrbracket \rrbracket = s'_2$ where $s'_1, s'_2 \in \Lambda_a^S$ in sharing normal form. Then thanks to Lemma 5 we know $s'_1 \rightarrow_{(\beta, D, L, C)}^* \llbracket M \rrbracket$ and $s'_2 \rightarrow_{(\beta, D, L, C)}^* \llbracket M \rrbracket$. Combined, we get confluence. $\square$

6 Conclusion, related work, and future directions

We have studied the interaction between the switch and the medial rule, the two characteristic inference rules of deep inference. We built a Curry–Howard interpretation based on this interaction, whose resulting calculus not only has the ability to duplicate terms atomically but can also duplicate solely the spine of an abstraction such that beta reduction can proceed on the duplicates. We show that this calculus has natural properties with respect to the λ -calculus.

This work, which started as an investigation into the Curry–Howard correspondence of the switch rule [25], fits into a broader effort to give a computational interpretation to intuitionistic deep-inference proof theory. Brännler and McKinley [9] give a natural reduction mechanism without medial (or switch), and observe that preservation of strong normalization fails. Guenot and Straßburger [14] investigate a different switch rule, corresponding to the implication-left rule of sequent calculus. He [17] extends the atomic λ -calculus to the $\lambda\mu$ -calculus.

Our future goal is to develop the intuitionistic open deduction formalism towards optimal reduction [23, 21, 3], via the remaining medial and switch rules [26].

Acknowledgements We thank the anonymous reviewers for their comments.

References

1. Abadi, M., Cardelli, L., Curien, P.L., Lévy, J.J.: Explicit substitutions. *Journal of Functional Programming* **1**(4), 375–416 (1991)
2. Accattoli, B., Kesner, D.: Preservation of strong normalisation modulo permutations for the structural lambda-calculus. *Logical Methods in Computer Science* **8**(1) (2012)
3. Asperti, A., Guerrini, S.: *The Optimal Implementation of Functional Programming Languages*. Cambridge University Press (1998)
4. Balabonski, T.: A unified approach to fully lazy sharing. *ACM SIGPLAN Notices* **47**(1), 469–480 (2012)
5. Balabonski, T.: Weak Optimality, and the Meaning of Sharing. In: *International Conference on Functional Programming (ICFP)*. pp. 263–274. Boston, United States (Sep 2013). <https://doi.org/10.1145/2500365.2500606>, <https://hal.archives-ouvertes.fr/hal-00907056>
6. Barendregt, H.P.: *The Lambda Calculus – Its Syntax and Semantics*, Studies in Logic and the Foundations of Mathematics, vol. 103. North-Holland (1984)
7. Berkling, K.J., Fehr, E.: A consistent extension of the lambda-calculus as a base for functional programming languages. *Information and Control* **55**, 89–101 (1982)
8. Blanc, T., Lévy, J.J., Maranget, L.: Sharing in the weak lambda-calculus. *Processes, Terms and Cycles: Steps on the Road to Infinity: Essays Dedicated to Jan Willem Klop on the Occasion of His 60th Birthday* **3838**, 70 (2005)
9. Brünnler, K., McKinley, R.: An algorithmic interpretation of a deep inference system. In: *International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR)*. pp. 482–496 (2008)
10. Brünnler, K., Tiu, A.: A local system for classical logic. In: *8th International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR)*. LNCS, vol. 2250, pp. 347–361 (2001)
11. Church, A., Rosser, J.B.: Some properties of conversion. *Transactions of the American Mathematical Society* **39**(3), 472–482 (1936), <http://www.jstor.org/stable/1989762>
12. Cockett, R., Seely, R.: Weakly distributive categories. *Journal of Pure and Applied Algebra* **114**(2), 133–173 (1997)
13. Fernández, M., Mackie, I., Sinot, F.R.: Lambda-calculus with director strings. *Applicable Algebra in Engineering, Communication and Computing* **15**(6), 393–437 (2005)
14. Guenot, N., Straßburger, L.: Symmetric normalisation for intuitionistic logic. In: *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* (2014)
15. Guglielmi, A., Gundersen, T., Parigot, M.: A proof calculus which reduces syntactic bureaucracy. In: *21st International Conference on Rewriting Techniques and Applications (RTA)*. pp. 135–150 (2010)
16. Gundersen, T., Heijltjes, W., Parigot, M.: Atomic lambda-calculus: a typed lambda-calculus with explicit sharing. In: *28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. pp. 311–320 (2013)
17. He, F.: *The Atomic Lambda-Mu Calculus*. Ph.D. thesis, University of Bath (2018)
18. Hendriks, D., van Oostrom, V.: Adbm. In: *19th International Conference on Automated Deduction (CADE)*. LNCS, vol. 2741, pp. 136–150 (2003)

19. Kennaway, R., Sleep, R.: Director strings as combinators. *ACM Transactions on Programming Languages and Systems* (1988)
20. Klop, J.W.: *Combinatory Reduction Systems*. Ph.D. thesis, Utrecht University (1980)
21. Lamping, J.: An algorithm for optimal lambda calculus reduction. In: *Proceedings of the 17th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. pp. 16–30 (1990)
22. Launchbury, J.: A natural semantics for lazy evaluation. In: *20th ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL)*. pp. 144–154 (1993)
23. Lévy, J.J.: *Optimal reductions in the lambda-calculus*. In: *To H.B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism*. Academic Press (1980)
24. van Oostrom, V., van de Looij, K.J., Zwitserlood, M.: Lambdascope: another optimal implementation of the lambda-calculus. In: *Workshop on Algebra and Logic on Programming Systems (ALPS)* (2004)
25. Sherratt, D.R.: *A lambda-calculus that achieves full laziness with spine duplication*. Ph.D. thesis, University of Bath (2019)
26. Tiu, A.: A local system for intuitionistic logic. In: *International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR)*. pp. 242–256 (2006)
27. Wadsworth, C.P.: *Semantics and Pragmatics of the Lambda-Calculus*. Ph.D. thesis, University of Oxford (1971)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.