Non-Deterministic Abstract Machines

Malgorzata Biernacka, Dariusz Biernacki, Serguei Lenglet, Alan Schmitt

Non-Deterministic Abstract Machines

Małgorzata Biernacka
Uniwersytet Wrocławski
Wrocław, Poland

Dariusz Biernacki
Uniwersytet Wrocławski
Wrocław, Poland

Sergueï Lenglet
Université de Lorraine
Nancy, France

Alan Schmitt
Inria
Rennes, France

Abstract

We present a generic design of abstract machines for non-deterministic programming languages, such as process calculi or concurrent lambda calculi, that provides a simple way to implement them. Such a machine traverses a term in the search for a redex, making non-deterministic choices when several paths are possible and backtracking when it reaches a dead end, i.e., an irreducible subterm. The search is guaranteed to terminate thanks to term annotations the machine introduces along the way.

We show how to automatically derive a non-deterministic abstract machine from a zipper semantics—a form of structural operational semantics in which the decomposition process of a term into a context and a redex is made explicit. The derivation method ensures the soundness and completeness of the machines w.r.t. the zipper semantics.

Keywords: Abstract machines, non-determinism

ACM Reference Format:
Małgorzata Biernacka, Dariusz Biernacki, Sergueï Lenglet, and Alan Schmitt. 2022. Non-Deterministic Abstract Machines. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Abstract machines, i.e., first-order tail-recursive transition systems for term reduction, such as SECD [27], CEK [11], and the KAM [26], are a traditional and celebrated artifact in the area of programming languages based on the λ-calculus. They serve both as a form of operational semantics [10, 11, 27] and an implementation model [24, 31] of programming languages, but they also play a role in other areas of logic in computer science, e.g., in proof theory [26] or higher-order model checking [39]. They are used as an implementation model also in concurrent languages [14, 16, 32, 35, 44], in particular to study distribution [4, 17–19, 22, 36].

Since in general designing a new abstract machine is a serious undertaking, several frameworks supporting mechanical or even automatic derivations of abstract machines from other forms of semantics have been developed [2, 5, 21, 42]. However, these frameworks assume a language that satisfies the unique decomposition property [5, 9], which entails that at each step one specific redex is selected, and thus the language follows a deterministic reduction strategy. This property does not hold in non-deterministic languages such as process calculi (or even in the λ-calculus without a fixed reduction strategy) and the existing methodology cannot be applied. Existing machines for non-deterministic languages are ad-hoc and may not be complete, i.e., not all reduction paths of the language can be simulated by the corresponding abstract machine [14, 16, 18, 32, 44].

This work presents a generic framework for the definition of complete abstract machines that implement the non-deterministic reduction relation in a systematic and controlled way. The idea is to go through a term to find a redex without following a specific strategy, picking arbitrarily a subterm when several are available—e.g., going left or right of an application in λ-calculus. The two main design principles are: (1) the machine should not remain stuck when it chooses a subterm which cannot reduce—in such a case we make it backtrack to its last choice; (2) the machine should not endlessly loop searching for redexes in subterms which cannot reduce—the machine annotates the subterms which are normal forms to prevent it from visiting them again.

Non-deterministic machines designed in this way can be complex even for small languages, therefore we show how to generate them automatically from an intermediary *zipper semantics*. This semantics, inspired by Huet [23], is a form of structural operational semantics (SOS) [38] that remembers the current position in a term by building a context, i.e., a syntactic object that represents a term with a hole [12]. This format of semantics makes it explicit how a term is decomposed into a context and a redex, and thus it can be seen as a non-deterministic counterpart of the decomposition function in (deterministic) context-based reduction semantics [8, 13]. While deterministic reduction semantics is directly implementable and the corresponding abstract machine can be viewed (roughly) as its optimization [9], non-deterministic reduction semantics, even when expressed as a zipper semantics, requires non-trivial instrumentation to become implementable in a complete way. Deriving the non-deterministic abstract machine (NDAM) from the zipper semantics consists exactly in such an instrumentation with the backtracking mechanism and normal-form annotations.

Conference'17, July 2017, Washington, DC, USA
2022. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

$$\text{init}\ \frac{}{t \xrightarrow[\text{app}]{\bullet} t'}\qquad \text{appL}\ \frac{t \xrightarrow[\text{app}]{@\,s :: \mathbb{E}} t'}{t \,@\, s \xrightarrow[\text{app}]{\mathbb{E}} t'}\qquad \text{appR}\ \frac{s \xrightarrow[\text{app}]{t\,@ :: \mathbb{E}} s'}{t \,@\, s \xrightarrow[\text{app}]{\mathbb{E}} t \,@\, s'}$$

$$\text{app}\lambda\ \frac{t \xrightarrow[\text{app}]{\lambda x :: \mathbb{E}} t'}{\lambda x.t \xrightarrow[\text{app}]{\mathbb{E}} t'}\qquad \text{app}\beta\ \frac{t \xrightarrow[\text{lam}]{s,\mathbb{E}} t'}{t \,@\, s \xrightarrow[\text{app}]{\mathbb{E}} t'}\qquad \text{lam}\beta\ \frac{}{\lambda x.t \xrightarrow[\text{lam}]{s,\mathbb{E}} \mathbb{E}[t\{s/x\}]}$$

Figure 1. Zipper semantics for the λ-calculus

We show how to derive an NDAM from an arbitrary zipper semantics that satisfies minimal conditions, and we prove that the resulting NDAM is sound and complete w.r.t. the semantics. Our approach applies in particular to process calculi, for which the abstract machines defined so far were ad-hoc and usually not complete.

The contributions of this paper are (1) a generic design of sound and complete, non-deterministic abstract machines which cannot get stuck or infinitely loop in a redex search, (2) with a systematic derivation procedure from an intermediary format, called zipper semantics. The resulting machine is an implementation of the non-deterministic source language.

We illustrate our method on the λ-calculus without a fixed reduction strategy and on a minimal process calculus HOcore [29], respectively in Sections 2 and 3. In Section 4 we consider a more complex calculus HOπ [40] that extends HOcore with name restriction, and we discuss the issues we need to address in this setting. We then give a derivation procedure of an NDAM from an arbitrary zipper semantics in Section 5 and we prove their correspondence in Section 6. We discuss related work in Section 7 and future work in Section 8. The appendix contains the proofs missing from the body of the paper. An implementation of the derivation procedure is also available [1].

2 Lambda-calculus

As a warm-up example, we present the zipper semantics and the corresponding NDAM for the λ-calculus with no fixed reduction strategy.

2.1 Syntax and Context-based Reduction Semantics

We let t, s range over λ-terms. We denote application with an explicit operator @ to annotate it later on. We represent a context $\mathbb{E}$ as a list of elementary contexts called *frames* $\mathfrak{F}$.

$$t, s ::= x \mid \lambda x.t \mid t \,@\, s$$
$$\mathfrak{F} ::= \lambda x \mid @\, t \mid t \,@ \qquad \mathbb{E}, \mathbb{F}, \mathbb{G} ::= \bullet \mid \mathfrak{F} :: \mathbb{E}$$

Because it is more convenient for the definition of the machine, we interpret contexts inside-out [10]: the head of the context is the innermost frame. The definition of plugging a term in a context $\mathbb{E}[t]$ is therefore as follows:

$$\bullet[t] \triangleq t \qquad (@\, s :: \mathbb{E})[t] \triangleq \mathbb{E}[t \,@\, s]$$
$$(\lambda x :: \mathbb{E})[t] \triangleq \mathbb{E}[\lambda x.t] \qquad (s \,@ :: \mathbb{E})[t] \triangleq \mathbb{E}[s \,@\, t]$$

We write $t\{s/x\}$ for the capture-avoiding substitution of x by s in t, and define the context-based reduction semantics $\to_{\text{rs}}$ of the λ-calculus by the following rule

$$\mathbb{E}[(\lambda x.t) \,@\, s] \to_{\text{rs}} \mathbb{E}[t\{s/x\}]$$

which can be read declaratively: if we find a redex in a context $\mathbb{E}$ built according to the given grammar of contexts, then we can reduce. This format of semantics does not make it apparent how to *decompose* a term to find a redex. On the other hand, structural operational semantics offers another common semantic format that makes it more explicit how to navigate in a term to find a redex, but it does not store the traversed path.

2.2 Zipper Semantics

A first step towards an abstract machine is to make explicit the step-by-step decomposition of a term into a context and a redex. To this end, we propose zipper semantics, a combination of SOS and reduction semantics. Like a regular SOS, a zipper semantics goes through a term looking for a redex using structural rules, except the current position in the term is made explicit with a context as in reduction semantics.

The zipper semantics for the λ-calculus is defined in Figure 1. It looks for a β-redex while constructing the surrounding context $\mathbb{E}$ at the same time. The decomposition happens in the rules appL, appR, and appλ, where we search for a redex by descending into the appropriate subterm of a given term. Each of these rules corresponds to a frame, with init initiating the search by setting the context to $\bullet$.

These rules actually look for the application at the root of the β-redex; checking that an application $t \,@\, s$ is indeed a β-redex is done by the rule appβ. It relies on an auxiliary transition $t \xrightarrow[\text{lam}]{s,\mathbb{E}} t'$, which checks that its source is indeed a λ-abstraction. In that case, we can β-reduce with rule lamβ. We can see that computation only occurs in the axiom; the other rules are simply propagating the result unchanged.

One may wonder why we need the rules appβ and lamβ while a single axiom $(\lambda x.t) \,@\, s \xrightarrow[\text{app}]{\mathbb{E}} \mathbb{E}[t\{s/x\}]$ is enough to recognize a β-redex. The reason is that we restrict ourselves to patterns discriminating only the head constructor of a term, to remain close to an abstract machine where the decomposition of a term occurs only one operator at a time.

Example 1. To illustrate further how to recognize a redex one operator at a time, suppose we restrict the argument of

the β-redex to a value $v ::= x \mid \lambda x.t$.

$$\mathbb{E}[(\lambda x.t)\ @\ v] \to_{\mathsf{rs}} \mathbb{E}[t\{v/x\}]$$

In such a case, we would need an extra transition $s \xrightarrow{x,t,\mathbb{E}}_{\mathsf{v}} t'$, checking that s is a value. The rule lamβ of Figure 1 would be replaced by the rule lamβ^{v} below.

$$\frac{s \xrightarrow{x,t,\mathbb{E}}_{\mathsf{v}} t'}{\lambda x.t \xrightarrow{s,\mathbb{E}}_{\mathsf{lam}} t'}\ \mathsf{lam}\beta^{\mathsf{v}} \qquad \frac{}{y \xrightarrow{x,t,\mathbb{E}}_{\mathsf{v}} \mathbb{E}[t\{y/x\}]}\ \mathsf{var}^{\mathsf{v}}$$

We omit the λ-abstraction rule lam$^{\mathsf{v}}$. □

The zipper semantics is sound and complete w.r.t. reduction semantics (cf Appendix A).

Theorem 2.1. *For all t and t', $t \to_{\mathsf{rs}} t'$ iff $t \to_{\mathsf{zs}} t'$.*

2.3 Non-Deterministic Abstract Machine

A zipper semantics derivation describes how to decompose a term into a redex and a context, but it is not yet an implementation, as it does not explain what to do when several rules can be applied, like appL, appR, and appβ in the application case. Our non-deterministic machine simply picks one of the rules to apply, and backtracks if it reaches a dead-end. We first present how we implement this backtracking and how it can be derived from the zipper semantics rules, before giving the formal definition of the NDAM.

Design principles. The decomposition at work in the zipper semantics rules can be directly turned into machine steps: looking at the source term between the conclusion and the premise, we see the change of focus occurring in the term. We introduce a machine mode for each transition kind (here, app and lam), and the rules appL, appR, and appβ are translated to the following *forward* machine steps, with $\mid$ separating the term from the context:

$$\langle t\ @\ s \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t \mid @\ s :: \mathbb{E}\rangle_{\mathsf{app}} \qquad \langle t\ @\ s \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t \mid s, \mathbb{E}\rangle_{\mathsf{lam}}$$
$$\langle t\ @\ s \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle s \mid t\ @ :: \mathbb{E}\rangle_{\mathsf{app}}$$

We see why interpreting the context inside-out is convenient: focusing on t in $\mathbb{E}[t\ @\ s]$ amounts to pushing the frame $@\ s$ on top of $\mathbb{E}$. It is the same as decomposing the term as $(@\ s :: \mathbb{E})[t]$: the innermost constructor becomes the topmost one in the context.

The resulting machine is non-deterministic as three different steps can be taken from the configuration $\langle t\ @\ s \mid \mathbb{E}\rangle_{\mathsf{app}}$. Unlike typical deterministic machines, it does not implement a particular strategy and does not choose, e.g., to always go left of an application as in the KAM [26]. A consequence is that the machine can make a wrong choice, i.e., focus on a term which cannot reduce, like a variable. In such cases, we want the machine to backtrack to the last configuration for which a choice had to be made, and no further.

To make backtracking possible, we record the applied rules in a stack π. When we reach a term which cannot reduce, we switch to a backtracking mode (here, bapp) where we can "unapply" a rule.

$$\langle t\ @\ s\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t\ ;\mathsf{appL} :: \pi \mid @\ s :: \mathbb{E}\rangle_{\mathsf{app}}$$
$$\langle x\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle \pi\ ;x \mid \mathbb{E}\rangle_{\mathsf{bapp}}$$
$$\langle \mathsf{appL} :: \pi\ ;t \mid @\ s :: \mathbb{E}\rangle_{\mathsf{bapp}} \mapsto \langle t\ @\ s\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}}$$

By doing so, the machine may try other rules on $t\ @\ s$, e.g., to find a redex in s. However, it should not try appL again, as the backtracking step implies there is no redex in t. We refer to backtracking steps like the last one as *backward*, and to steps like the middle one as *switching*. We see that the backward step is simply the reverse of the corresponding forward step.

We prevent the machine from choosing a previously explored path by annotating the root operator of an already tested subterm. An annotation $t\ @^{\mathsf{app}}\ s$ means that $t\ @\ s$ has already been tried for $\xrightarrow{\mathbb{E}}_{\mathsf{app}}$ transitions and is a normal form for it. Similarly, a term annotated lam is a normal form w.r.t. $\xrightarrow{s,\mathbb{E}}_{\mathsf{lam}}$ (it is not a λ-abstraction). A term can be annotated with both app and lam, for instance if it is a variable.

The machine can take a forward step only if the term in focus has not been already tested. For $t\ @\ s$, we can try appL (resp. appR) only if t (resp. s) is not annotated with app, and appβ only if t is not annotated with lam. If none of the steps applies because of the annotations, then all possible rules have been tried and $t\ @\ s$ is a normal form for app: the machine should backtrack and annotate the term accordingly. In what follows, Σ represents an annotation set.

$$\langle x^{\Sigma}\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle \pi\ ;x^{\Sigma\cup\{\mathsf{app}\}} \mid \mathbb{E}\rangle_{\mathsf{bapp}}$$
$$\langle t\ @^{\Sigma}\ s\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle \pi\ ;t\ @^{\Sigma\cup\{\mathsf{app}\}}\ s \mid \mathbb{E}\rangle_{\mathsf{bapp}} \quad \text{if no other step applies}$$

We see that switching steps are of two kinds: either the language construct does not have a forward step for a given mode (like a variable in the app mode), or all possible rules have been tried for the construct. They both can be derived from the zipper semantics by looking at which rule can be applied to each construct. This derivation is made easier by the constraint that the decomposition occurs one operator at a time in zipper rules. If we allowed for more complex patterns such as $(\lambda x.t)\ @\ s$, we would have to create a switching step for the terms not fitting this pattern, like $x\ @\ s$, and enumerating these anti-patterns would be more difficult [25].

Finally, because we store the annotations of a term in its root operator, we need to remember them when a forward step removes the operator, to be able to restore them when we backtrack. We do so in the stack π.

$$\langle t\ @^{\Sigma}\ s\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t\ ;(\mathsf{appL}, \Sigma) :: \pi \mid @\ s :: \mathbb{E}\rangle_{\mathsf{app}}$$
$$\langle (\mathsf{appL}, \Sigma) :: \pi\ ;t \mid @\ s :: \mathbb{E}\rangle_{\mathsf{bapp}} \mapsto \langle t\ @^{\Sigma}\ s\ ;\pi \mid \mathbb{E}\rangle_{\mathsf{app}}$$

In this simple example we could do without the stack because the contexts encode precisely the rules that have been

applied along the way. In general, however, a single context cannot always reflect the derivation tree, as we can see in the HOπ example (Section 4).

The next example illustrates how annotations work, and also that they may no longer hold after reduction. Therefore they should be erased before searching for the next redex.

Example 2. Let $\Omega \triangleq (\lambda^{\emptyset}x.x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset}) \,@^{\emptyset} (\lambda^{\emptyset}x.x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset})$. We show a possible machine run for this term, where we label forward and backward steps with the rule they apply or unapply, and switching steps with a constant τ. For readability, we write only the term under focus.

The machine first goes left and under the λ-abstraction.

$$\langle \Omega \mid \ldots \rangle_{\mathsf{app}} \overset{\mathsf{appL}}{\longmapsto} \overset{\mathsf{app}\lambda}{\longmapsto} \langle x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset} \mid \ldots \rangle_{\mathsf{app}}$$

At that point, it may test whether the application is a β-redex. Since it is not the case, it backtracks, annotating the variable in function position.

$$\langle x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset} \mid \ldots \rangle_{\mathsf{app}} \overset{\mathsf{app}\beta}{\longmapsto} \langle x^{\emptyset} \mid \ldots \rangle_{\mathsf{lam}} \overset{\tau}{\mapsto} \overset{-\mathsf{app}\beta}{\longmapsto} \langle x^{\mathsf{lam}} \,@^{\emptyset}\, x^{\emptyset} \mid \ldots \rangle_{\mathsf{app}}$$

From there, it necessarily tests the other possibilities appL and appR (in no predefined order), and fails in both cases.

$$\langle x^{\mathsf{lam}} \,@^{\emptyset}\, x^{\emptyset} \mid \ldots \rangle_{\mathsf{app}} \overset{\mathsf{appL}}{\longmapsto} \overset{\tau}{\mapsto} \overset{-\mathsf{appL}}{\longmapsto} \overset{\mathsf{appR}}{\longmapsto} \overset{\tau}{\mapsto} \overset{-\mathsf{appR}}{\longmapsto} \langle x^{\{\mathsf{app},\mathsf{lam}\}} @^{\emptyset} x^{\mathsf{app}} \mid \ldots \rangle_{\mathsf{app}}$$

Then it can only backtrack to reconstruct the λ-abstraction on the left, and then the whole term.

$$\langle x^{\{\mathsf{app},\mathsf{lam}\}} @^{\emptyset} x^{\mathsf{app}} \mid \ldots \rangle_{\mathsf{app}} \overset{\tau}{\mapsto} \overset{-\mathsf{app}\lambda}{\longmapsto}$$

$$\langle \lambda^{\emptyset}x.x^{\{\mathsf{app},\mathsf{lam}\}} @^{\mathsf{app}} x^{\mathsf{app}} \mid \ldots \rangle_{\mathsf{app}} \overset{\tau}{\mapsto} \overset{-\mathsf{appL}}{\longmapsto}$$

$$\langle (\lambda^{\mathsf{app}}x.x^{\{\mathsf{app},\mathsf{lam}\}} @^{\mathsf{app}} x^{\mathsf{app}}) \,@^{\emptyset} (\lambda^{\emptyset}x.x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset}) \mid \ldots \rangle_{\mathsf{app}}$$

The machine can then look for a redex in the λ-abstraction on the right, and it would result in the same annotations as for the one on the left, not necessarily generated in the same order. It can also rightfully recognize the term as a β-redex, with the sequence $\overset{\mathsf{app}\beta}{\longmapsto} \overset{\mathsf{lam}\beta}{\longmapsto}$, the last step performing the reduction. After the reduction, we should also erase the remaining annotations. If we do not erase them, the result of the reduction would be

$$\langle (\lambda^{\emptyset}x.x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset}) @^{\{\mathsf{app}\}} (\lambda^{\emptyset}x.x^{\emptyset} \,@^{\emptyset}\, x^{\emptyset}) \mid \ldots \rangle_{\mathsf{app}}$$

and the app annotation would wrongfully signal the term as a normal-form, preventing it from being reduced. □

Formal definition. We let α range over annotations, Σ over annotation sets, and denote the empty set by $\emptyset$. We extend the λ-calculus syntax as follows:

$$\alpha ::= \mathsf{app} \mid \mathsf{lam} \qquad t, s ::= x^{\Sigma} \mid \lambda^{\Sigma}x.t \mid t \,@^{\Sigma}\, s$$

We write $\mathsf{an}(t)$ for the annotation set at the root of t, e.g., $\mathsf{an}(t \,@^{\Sigma}\, s) \triangleq \Sigma$. We write $t^{\cup\alpha}$ for its extension with α so that $\mathsf{an}(t^{\cup\alpha}) = \mathsf{an}(t) \cup \{\alpha\}$. We write $|t|$ for the erasure of t, where all the annotation sets in t are made empty.

The syntax of contexts now uses annotated terms, and plugging returns an annotated term such that the annotation sets of the context operators are chosen empty: e.g., $(\lambda x :: \mathbb{E})[t] \triangleq \mathbb{E}[\lambda^{\emptyset}x.t]$. Plugging is used only after a reduction step, where all the annotation sets are erased anyway.

We let ρ range over rule names and π over rule stacks, defined as $\pi ::= \mathsf{init} \mid (\rho, \Sigma) :: \pi$. The definition of the machine for the λ-calculus is given in Figure 2.

A *forward* configuration $\langle t \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{m}}$ (with $\mathsf{m} \in \{\mathsf{app}, \mathsf{lam}\}$) discriminates on (the root operator of) t to apply a rule of the zipper semantics. For an inductive rule, it results in a change of focus and an extension of the stack, on which we record the applied rule and the annotation set of the root operator. Taking such a step is possible only if the new term under focus is not a normal form. A special case of forward step is the *initial* one from $\langle t \rangle_{\mathsf{zs}}$ which does not have a side-condition, as we assume the annotation sets of t to be empty.

The axiom (lamβ) is implemented by the first transition for the lam machine mode. The redex has been found and backtracking is no longer necessary: we drop the stack, reconstruct the entire term, and switch to the initial mode to search for a new redex starting from the root of the term. We also erase all annotations, as they may no longer be valid in the new term, as illustrated in Example 2.

If a forward configuration cannot apply a rule, we switch to the corresponding *backward* mode, annotating t in the process: these are the two "otherwise" steps. A backward configuration $\langle \pi \,;\, t \mid \mathbb{E} \rangle_{\mathsf{bm}}$ inspects the stack π to unapply the rule at its top. While a backward step restores the configuration of the corresponding forward step, the term contains more annotations after a backward step than before taking the forward step: in $\langle \pi \,;\, t \mid \mathbb{E} \rangle_{\mathsf{bm}}$, we have $\mathsf{m} \in \mathsf{an}(t)$ by construction. The annotations prevent the machine from reapplying a rule it just unapplied. The *normal form* mode $\langle t \rangle_{\mathsf{nf}}$ signals that the term cannot reduce.

A machine run starts with an initial configuration $\langle t \rangle_{\mathsf{zs}}$ where all the annotation sets of t are empty. The semantics of the machine is given by these configurations: if $\langle t \rangle_{\mathsf{zs}} \mapsto^{+} \langle t' \rangle_{\mathsf{zs}}$ such that the sequence $\mapsto^{+}$ does not go through another initial configuration, then $t \to_{\mathsf{zs}} t'$. Similarly, if $\langle t \rangle_{\mathsf{zs}} \mapsto^{+} \langle t' \rangle_{\mathsf{nf}}$, then $|t'| = t$ and t is a normal form. We prove the correspondence and termination independently from the source zipper semantics in Section 6.

3 HOcore

We consider a minimal process calculus called HOcore [29], which can be seen as an extension of the λ-calculus with parallel composition.

$$\langle t\rangle_{\mathsf{zs}} \mapsto \langle t\,;\mathsf{init} \mid \bullet\rangle_{\mathsf{app}} \qquad \langle \mathsf{init}\,; t \mid \bullet\rangle_{\mathsf{bapp}} \mapsto \langle t\rangle_{\mathsf{nf}}$$

$$\begin{array}{lll}
\langle t\ @^{\Sigma}\ s\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t\,;(\mathsf{appL},\Sigma)::\pi \mid @\ s::\mathbb{E}\rangle_{\mathsf{app}} & \text{if app} \notin \mathsf{an}(t) & \langle (\mathsf{appL},\Sigma)::\pi\,; t \mid @\ s::\mathbb{E}\rangle_{\mathsf{bapp}} \mapsto \langle t\ @^{\Sigma}\ s\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \\
\langle t\ @^{\Sigma}\ s\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle s\,;(\mathsf{appR},\Sigma)::\pi \mid t\ @::\mathbb{E}\rangle_{\mathsf{app}} & \text{if app} \notin \mathsf{an}(s) & \langle (\mathsf{appR},\Sigma)::\pi\,; s \mid t\ @::\mathbb{E}\rangle_{\mathsf{bapp}} \mapsto \langle t\ @^{\Sigma}\ s\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \\
\langle t\ @^{\Sigma}\ s\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t\,;(\mathsf{app}\beta,\Sigma)::\pi \mid s,\mathbb{E}\rangle_{\mathsf{lam}} & \text{if lam} \notin \mathsf{an}(t) & \langle (\mathsf{app}\beta,\Sigma)::\pi\,; t \mid s,\mathbb{E}\rangle_{\mathsf{blam}} \mapsto \langle t\ @^{\Sigma}\ s\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \\
\langle \lambda^{\Sigma} x.t\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle t\,;(\mathsf{app}\lambda,\Sigma)::\pi \mid \lambda x::\mathbb{E}\rangle_{\mathsf{app}} & \text{if app} \notin \mathsf{an}(t) & \langle (\mathsf{app}\lambda,\Sigma)::\pi\,; t \mid \lambda x::\mathbb{E}\rangle_{\mathsf{bapp}} \mapsto \langle \lambda^{\Sigma} x.t\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \\
\langle t\,;\pi \mid \mathbb{E}\rangle_{\mathsf{app}} \mapsto \langle \pi\,; t^{\cup\mathsf{app}} \mid \mathbb{E}\rangle_{\mathsf{bapp}} & \text{otherwise} &
\end{array}$$

$$\begin{array}{lll}
\langle \lambda^{\Sigma} x.t\,;\pi \mid s,\mathbb{E}\rangle_{\mathsf{lam}} \mapsto \langle |\mathbb{E}[t\{s/x\}]|\rangle_{\mathsf{zs}} & & \\
\langle t\,;\pi \mid s,\mathbb{E}\rangle_{\mathsf{lam}} \mapsto \langle \pi\,; t^{\cup\mathsf{lam}} \mid s,\mathbb{E}\rangle_{\mathsf{blam}} & \text{otherwise} &
\end{array}$$

Figure 2. Non-Deterministic Abstract Machine for the λ-calculus

$$\mathsf{init}\ \frac{P \xrightarrow{\bullet}_{\mathsf{par}} P'}{P \to_{\mathsf{zs}} P'} \qquad \mathsf{parL}\ \frac{P \xrightarrow{\| Q::\mathbb{E}}_{\mathsf{par}} P'}{P \| Q \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'}\ (s) \qquad \mathsf{parOutL}\ \frac{P \xrightarrow{\bullet,\mathcal{L},\mathbb{E},Q}_{\mathsf{out}} P'}{P \| Q \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'} \qquad \mathsf{parOutR}\ \frac{Q \xrightarrow{\bullet,\mathcal{R},\mathbb{E},P}_{\mathsf{out}} P'}{P \| Q \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'} \qquad \mathsf{outParL}\ \frac{P \xrightarrow{\| Q::\mathbb{F},\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'}{P \| Q \xrightarrow{\mathbb{F},\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'}\ (s)$$

$$\mathsf{outIn}\ \frac{R \xrightarrow{\bullet,\mathcal{S},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} P'}{\overline{a}\langle P\rangle \xrightarrow{\mathbb{F},\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'} \qquad \mathsf{inParL}\ \frac{R \xrightarrow{\| Q::\mathbb{G},\mathcal{S},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} P'}{R \| Q \xrightarrow{\mathbb{G},\mathcal{S},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} P'}\ (s) \qquad \mathsf{inComL}\ \frac{b = a}{b(X).R \xrightarrow{\mathbb{G},\mathcal{L},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} \mathbb{E}[\mathbb{F}[\mathbf{0}] \| \mathbb{G}[R\{P/X\}]]}\ (s)$$

Figure 3. Output-first Zipper Semantics for HOcore

3.1 Syntax and Semantics

We let a, b range over *channel names*, X, Y over *process variables*, and we define the syntax of processes as follows.

$$P, Q, R ::= X \mid \mathbf{0} \mid P \| Q \mid a(X).P \mid \overline{a}\langle P\rangle$$

The process $\mathbf{0}$ is the inactive process, $P \| Q$ runs P and Q in parallel, and a communication may happen between an input $a(X).P$ and an output $\overline{a}\langle Q\rangle$ that run in parallel. The communication is asynchronous because a message output does not have a continuation [41]; we discuss the synchronous case in Remark 1. In spite of its minimal number of constructors, HOcore is Turing-complete [29].

The semantics of process calculi is usually presented either with a structural congruence relation which reorders terms to make redexes appear, bringing input and output processes together, or with a labeled transition system which preserves the structure of the term [41]. Instead, we present it first as a reduction semantics with explicit contexts, as in Section 2.1, which makes it easier to come up with (or translate into) the corresponding zipper semantics.

We define frames as $\mathfrak{F} ::= \| P \mid P \|$ and plugging as follows.

$$\bullet[P] \triangleq P \quad (\| Q::\mathbb{E})[P] \triangleq \mathbb{E}[P \| Q] \quad (Q \| ::\mathbb{E})[P] \triangleq \mathbb{E}[Q \| P]$$

A redex is a parallel composition with an input on one side and an output on the same name on the other side, both surrounded with contexts. The general formulation of such communication sites in a program can be expressed with the following reduction semantics, where we write $P\{Q/X\}$ for the capture-avoiding substitution of X by Q in P:

$$\mathbb{E}[\mathbb{F}[\overline{a}\langle Q\rangle] \| \mathbb{G}[a(X).P]] \to_{\mathsf{rs}} \mathbb{E}[\mathbb{F}[\mathbf{0}] \| \mathbb{G}[P\{Q/X\}]]$$
$$\mathbb{E}[\mathbb{G}[a(X).P] \| \mathbb{F}[\overline{a}\langle Q\rangle]] \to_{\mathsf{rs}} \mathbb{E}[\mathbb{G}[P\{Q/X\}] \| \mathbb{F}[\mathbf{0}]]$$

3.2 Zipper Semantics

Finding a HOcore redex requires us to recognize three constructs (parallel composition along with output and input on a shared name) and build the contexts $\mathbb{E}$, $\mathbb{F}$, and $\mathbb{G}$. The first step is to find the parallel composition; once the communicating processes $P \| Q$ are found, the communication rules of typical LTSs for process calculi [29, 40, 41] have two premises looking for the output and the input in P and Q respectively. To be closer to an abstract machine, we sequentialize the search by looking for the output first (while constructing $\mathbb{F}$) and then the input (with $\mathbb{G}$)—the opposite choice would produce a completely symmetric semantics. Figure 3 presents such an output-first zipper semantics, where we omit the symmetric versions of the rules marked with the

symbol (s)—a convention we follow from now on. The resulting semantics is close to complementary semantics [30], where the communication is also sequentialized.

The transition $\xrightarrow{\mathbb{E}}_{\mathsf{par}}$ is looking for the parallel composition while building $\mathbb{E}$: it proceeds as $\xrightarrow{\mathbb{E}}_{\mathsf{app}}$ in the λ-calculus. Once we find the parallel composition, we look for the output either on the left or on the right with respectively rules parOutL and parOutR. We record the side we pick with a parameter $\mathcal{S} ::= \mathcal{L} \mid \mathcal{R}$. For example, in rule parOutL, we look for an output in P on the left ($\mathcal{L}$), remembering that we should later search for a corresponding input in Q. We also initialize the context $\mathbb{F}$ surrounding the output with $\bullet$ and remember $\mathbb{E}$ as the context enclosing the whole redex.

The transition $\xrightarrow{\mathbb{F},\mathcal{S},\mathbb{E},R}_{\mathsf{out}}$ decomposes its source process to find an output, building $\mathbb{F}$ at the same time: the other parameters $\mathcal{S}$, $\mathbb{E}$, and R remain unchanged during the search. When we find the output $\overline{a}\langle P\rangle$ (rule outIn), we look for a corresponding input in R using $\xrightarrow{\mathbb{G},\mathcal{S},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}}$, which builds the context $\mathbb{G}$ during the search. Once we find an input on a, we compute the result of the communication, which depends whether the output is on the left (rule inComL) or on the right (omitted rule inComR).

The two HOcore semantics coincide (Appendix B).

Theorem 3.1. *For all P and P', $P \to_{\mathsf{zs}} P'$ iff $P \to_{\mathsf{rs}} P'$.*

Remark 1 (Synchronous communication)**.** For a synchronous calculus with an output $\overline{a}\langle P\rangle Q$, the rule outIn would pass the continuation Q as an argument of the input transition in. The continuation Q would then be plugged into $\mathbb{F}$ in the axioms inComL and inComR.

Remark 2 (Left-first or right-first searches)**.** After finding the communicating processes $P \parallel Q$, we could always go left (in P). When we find an output or input in P, we look for its complement in Q. A right-first approach is also possible. We present the left-first zipper semantics and its corresponding machine in Appendix B. The left-first semantics corresponds to a depth-first traversal of the redexes of Section 3.1, but it requires four kinds of transition instead of three for the output-first one. More importantly, it does not scale to HOπ, as explained in Remark 3.

3.3 Non-Deterministic Abstract Machine

We derive the HOcore NDAM from its zipper semantics along the same principles as for the λ-calculus: each rule of the semantics corresponds to a forward step and a backward step, and when no forward step applies to a configuration, we switch to a backward configuration. The difference is in the normal-form annotations: in λ-calculus, to be a normal form w.r.t. $\xrightarrow{s,\mathbb{E}}_{\mathsf{lam}}$ or $\xrightarrow{\mathbb{E}}_{\mathsf{app}}$ does not depend on the arguments s and $\mathbb{E}$. In HOcore, being a normal form depends on some of the arguments in the input and output transitions.

For example, in a process $(\overline{a}\langle \mathbf{0}\rangle \parallel \overline{b}\langle \mathbf{0}\rangle) \parallel Q$, we may look into Q for an input on a or on b. If Q does not contain an input on a, then annotating it with the mode in would prevent from searching in Q for an input on b. We therefore include the name in the annotation, marking the root operator of Q with (in, a), meaning that Q cannot do an input on a. If it also cannot do an input on b, then its root operator will be annotated with both (in, a) and (in, b).

With outputs the problem is similar, but not completely symmetric. Let $P_{a,b} = \overline{a}\langle \mathbf{0}\rangle \parallel \overline{b}\langle \mathbf{0}\rangle$, and consider a process $(P_{a,b} \parallel Q) \parallel R$. We may try to find a communication between $P_{a,b}$ and Q first. If Q does not contain an input on a or b, then $P_{a,b}$ is a normal form w.r.t. the output search transition $\xrightarrow{\bullet,\mathcal{L},\parallel R ::\, \bullet, Q}_{\mathsf{out}}$, but a communication between $P_{a,b}$ and R is still possible. As a result, we annotate the root operator of $P_{a,b}$ with (out, Q), meaning that the outputs of $P_{a,b}$ are not complemented by the inputs in Q. Such an annotation does not prevent trying to make $P_{a,b}$ and R communicate, which would correspond to the transition $\xrightarrow{\parallel Q ::\, \bullet, \mathcal{L}, \bullet, R}_{\mathsf{out}}$.

As before, Σ ranges over annotation sets, and $|P|$ is the erasure of P, the annotated process with empty annotation sets. The syntax of annotations and processes is as follows.

$$\alpha ::= \mathsf{par} \mid (\mathsf{out}, |P|) \mid (\mathsf{in}, a)$$

$$P, Q, R ::= X^{\Sigma} \mid \mathbf{0}^{\Sigma} \mid P \parallel^{\Sigma} Q \mid a^{\Sigma}(X).P \mid \overline{a}^{\Sigma}\langle P\rangle$$

Substitution and plugging are extended to annotated processes as expected. The definition of the machine is given in Figure 4. The process P in an annotation $(\mathsf{out}, |P|)$—as in the side conditions in the par-transitions—is erased, because normal forms are defined with respect to the zipper semantics transitions, where processes are not annotated. Apart from richer annotations, the definition of the machine follows the principles of Section 2.3. Note that the "otherwise" step for the input mode includes the operators that are not parsed in that mode, but also the inputs on a name distinct from a.

4 HOπ

Finally, we present the zipper semantics of HOπ, an extension of HOcore with name restriction. The main difficulty is that the evaluation contexts surrounding the communicating processes can be themselves modified by the reduction.

4.1 Syntax and Semantics

We add name restriction to HOcore processes and frames.

$$P, Q, R ::= \ldots \mid \nu a.P \qquad \mathfrak{F} ::= \ldots \mid \nu a$$

To remain close to HOcore, the calculus of this section is asynchronous: outputs $\overline{a}\langle P\rangle$ do not have a continuation, unlike the original HOπ [40], Adding continuations would not be an issue as pointed out in Remark 1.

The scope of a in $\nu a.P$ is restricted to P, so that a communication on a is possible inside P only. For instance, the process $a(X).X \parallel \nu a.\overline{a}\langle \mathbf{0}\rangle$ cannot reduce, because the name a

$$
\begin{array}{rcll}
\langle P\rangle_{\mathsf{zs}} & \mapsto & \langle P\,;\mathsf{init}\mid\bullet\rangle_{\mathsf{par}} & \\
\langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & \mapsto & \langle P\,;(\mathsf{parL},\Sigma)::\pi\mid\parallel Q::\mathbb{E}\rangle_{\mathsf{par}} & \text{if par}\notin\mathsf{an}(P) \quad (s)\\
\langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & \mapsto & \langle P\,;(\mathsf{parOutL},\Sigma)::\pi\mid\bullet,\mathcal{L},\mathbb{E},Q\rangle_{\mathsf{out}} & \text{if }(\mathsf{out},|Q|)\notin\mathsf{an}(P)\\
\langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & \mapsto & \langle Q\,;(\mathsf{parOutR},\Sigma)::\pi\mid\bullet,\mathcal{R},\mathbb{E},P\rangle_{\mathsf{out}} & \text{if }(\mathsf{out},|P|)\notin\mathsf{an}(Q)\\
\langle P\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & \mapsto & \langle\pi\,;P^{\cup\mathsf{par}}\mid\mathbb{E}\rangle_{\mathsf{bpar}} & \text{otherwise}\\
\langle\mathsf{init}\,;P\mid\bullet\rangle_{\mathsf{bpar}} & \mapsto & \langle P\rangle_{\mathsf{nf}} & \\
\langle(\mathsf{parL},\Sigma)::\pi\,;P\mid\parallel Q::\mathbb{E}\rangle_{\mathsf{bpar}} & \mapsto & \langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & (s)\\
\langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{out}} & \mapsto & \langle P\,;(\mathsf{outParL},\Sigma)::\pi\mid\parallel Q::\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{out}} & \text{if }(\mathsf{out},|R|)\notin\mathsf{an}(P) \quad (s)\\
\langle\overline{a}^{\Sigma}\langle P\rangle\,;\pi\mid\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{out}} & \mapsto & \langle R\,;(\mathsf{outIn},\Sigma)::\pi\mid\bullet,\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & \text{if }(\mathsf{in},a)\notin\mathsf{an}(R)\\
\langle P\,;\pi\mid\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{out}} & \mapsto & \langle\pi\,;P^{\cup(\mathsf{out},|R|)}\mid\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{bout}} & \text{otherwise}\\
\langle(\mathsf{parOutL},\Sigma)::\pi\,;P\mid\bullet,\mathcal{L},\mathbb{E},Q\rangle_{\mathsf{bout}} & \mapsto & \langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & \\
\langle(\mathsf{parOutR},\Sigma)::\pi\,;Q\mid\bullet,\mathcal{R},\mathbb{E},P\rangle_{\mathsf{bout}} & \mapsto & \langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{E}\rangle_{\mathsf{par}} & \\
\langle(\mathsf{outParL},\Sigma)::\pi\,;P\mid\parallel Q::\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{bout}} & \mapsto & \langle P\parallel^{\Sigma}Q\,;\pi\mid\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{out}} & (s)\\
\langle R\parallel^{\Sigma}Q\,;\pi\mid\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & \mapsto & \langle R\,;(\mathsf{inParL},\Sigma)::\pi\mid\parallel Q::\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & \text{if }(\mathsf{in},a)\notin\mathsf{an}(R) \quad (s)\\
\langle b^{\Sigma}(X).R\,;\pi\mid\mathbb{G},\mathcal{L},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & \mapsto & \langle|\mathbb{E}[\mathbb{F}[\mathbf{0}]\parallel\mathbb{G}[R\{P/X\}]]|\rangle_{\mathsf{zs}} & \text{if } a=b\\
\langle b^{\Sigma}(X).R\,;\pi\mid\mathbb{G},\mathcal{R},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & \mapsto & \langle|\mathbb{E}[\mathbb{G}[R\{P/X\}]\parallel\mathbb{F}[\mathbf{0}]]|\rangle_{\mathsf{zs}} & \text{if } a=b\\
\langle R\,;\pi\mid\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & \mapsto & \langle\pi\,;R^{\cup(\mathsf{in},a)}\mid\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{bin}} & \text{otherwise}\\
\langle(\mathsf{outIn},\Sigma)::\pi\,;R\mid\bullet,\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{bin}} & \mapsto & \langle\overline{a}^{\Sigma}\langle P\rangle\,;\pi\mid\mathbb{F},\mathbb{S},\mathbb{E},R\rangle_{\mathsf{out}} & \\
\langle(\mathsf{inParL},\Sigma)::\pi\,;R\mid\parallel Q::\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{bin}} & \mapsto & \langle R\parallel^{\Sigma}Q\,;\pi\mid\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}\rangle_{\mathsf{in}} & (s)
\end{array}
$$

Figure 4. Non-Deterministic Abstract Machine for HOcore

$$
\mathsf{init}\ \frac{P\xrightarrow{\bullet}_{\mathsf{par}}P'}{P\rightarrow_{\mathsf{zs}}P'}
\qquad
\mathsf{parNu}\ \frac{P\xrightarrow{va::\mathbb{E}}_{\mathsf{par}}P'}{va.P\xrightarrow{\mathbb{E}}_{\mathsf{par}}P'}
\qquad
\mathsf{parL}\ \frac{P\xrightarrow{\parallel Q::\mathbb{E}}_{\mathsf{par}}P'}{P\parallel Q\xrightarrow{\mathbb{E}}_{\mathsf{par}}P'}\ (s)
$$

$$
\mathsf{parOutL}\ \frac{P\xrightarrow{\bullet,\bullet,\mathcal{L},\mathbb{E},Q}_{\mathsf{out}}P'}{P\parallel Q\xrightarrow{\mathbb{E}}_{\mathsf{par}}P'}
\qquad
\mathsf{parOutR}\ \frac{Q\xrightarrow{\bullet,\bullet,\mathcal{R},\mathbb{E},P}_{\mathsf{out}}P'}{P\parallel Q\xrightarrow{\mathbb{E}}_{\mathsf{par}}P'}
$$

$$
\mathsf{outParL}\ \frac{P\xrightarrow{\parallel Q::\mathbb{F}_1,\mathbb{F}_2,\mathbb{S},\mathbb{E},R}_{\mathsf{out}}P'}{P\parallel Q\xrightarrow{\mathbb{F}_1,\mathbb{F}_2,\mathbb{S},\mathbb{E},R}_{\mathsf{out}}P'}\ (s)
\qquad
\mathsf{outNu}\ \frac{P\xrightarrow{\mathbb{F}_1,vb::\mathbb{F}_2,\mathbb{S},\mathbb{E},R}_{\mathsf{out}}P'}{vb.P\xrightarrow{\mathbb{F}_1,\mathbb{F}_2,\mathbb{S},\mathbb{E},R}_{\mathsf{out}}P'}
$$

$$
\mathsf{outIn}\ \frac{R\xrightarrow{\bullet,\mathbb{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}}P'\qquad a\notin\mathsf{bn}(\mathbb{F}_2)}{\overline{a}\langle P\rangle\xrightarrow{\mathbb{F}_1,\mathbb{F}_2,\mathbb{S},\mathbb{E},R}_{\mathsf{out}}P'}
$$

$$
\mathsf{inParL}\ \frac{R\xrightarrow{\parallel Q::\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}}P'}{R\parallel Q\xrightarrow{\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}}P'}\ (s)
\qquad
\mathsf{inNu}\ \frac{R\xrightarrow{vb::\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}}P'\qquad a\neq b}{vb.R\xrightarrow{\mathbb{G},\mathbb{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}}P'}
$$

$$
\mathsf{inComL}\ \frac{a=b}{b(X).R\xrightarrow{\mathbb{G},\mathcal{L},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}}\mathbb{E}[\mathbb{F}_2[\mathbb{F}_1[\mathbf{0}]\parallel\mathbb{G}[R\{P/X\}]]]}\ (s)
$$

Figure 5. Zipper Semantics for HOπ

is restricted to the process on the right. In general, a process $\mathbb{E}[\overline{a}\langle P\rangle]$ or $\mathbb{E}[a(X).P]$ cannot communicate on a if $\mathbb{E}$ captures a. To check this, we compute the set of names bound by $\mathbb{E}$, written $\mathsf{bn}(\mathbb{E})$, as follows.

$$
\begin{array}{ll}
\mathsf{bn}(\bullet)\triangleq\emptyset & \mathsf{bn}(\parallel P::\mathbb{E})\triangleq\mathsf{bn}(\mathbb{E})\\
\mathsf{bn}(va::\mathbb{E})\triangleq\{a\}\cup\mathsf{bn}(\mathbb{E}) & \mathsf{bn}(P\parallel::\mathbb{E})\triangleq\mathsf{bn}(\mathbb{E})
\end{array}
$$

Name restriction does not forbid the communication on unrestricted names, but the scope of restricted names has to be enlarged to prevent them from escaping their delimiter. For example, we have

$$b(X).(X \parallel \overline{c}\langle \mathbf{0}\rangle) \parallel va.(\overline{b}\langle a(Y).Y\rangle \parallel \overline{a}\langle \mathbf{0}\rangle) \to_{\mathrm{rs}} va.(a(Y).Y \parallel \overline{c}\langle \mathbf{0}\rangle \parallel \mathbf{0} \parallel \overline{a}\langle \mathbf{0}\rangle)$$

The scope of a has been extended to include the receiving process on b. This phenomenon is known as *scope extrusion*. To reflect it at the level of contexts, we define an operation $\mathrm{extr}(\mathbb{E})$ which returns a pair of contexts $(\mathbb{E}_1, \mathbb{E}_2)$ such that $\mathbb{E}_2$ contains the binding frames, while $\mathbb{E}_1$ contains the remaining frames. We assume free names to be distinct from bound names using α-conversion if necessary, to avoid capture during extrusion.

$$\mathrm{extr}(\bullet) \triangleq (\bullet, \bullet) \qquad \frac{\mathrm{extr}(\mathbb{E}) = (\mathbb{E}_1, \mathbb{E}_2)}{\mathrm{extr}(va :: \mathbb{E}) \triangleq (\mathbb{E}_1, va :: \mathbb{E}_2)}$$

$$\frac{\mathrm{extr}(\mathbb{E}) = (\mathbb{E}_1, \mathbb{E}_2)}{\mathrm{extr}(\parallel P :: \mathbb{E}) \triangleq (\parallel P :: \mathbb{E}_1, \mathbb{E}_2)} \qquad \frac{\mathrm{extr}(\mathbb{E}) = (\mathbb{E}_1, \mathbb{E}_2)}{\mathrm{extr}(P \parallel :: \mathbb{E}) \triangleq (P \parallel :: \mathbb{E}_1, \mathbb{E}_2)}$$

We define the reduction semantics $\to_{\mathrm{rs}}$ of HOπ as follows, assuming $a \notin \mathrm{bn}(\mathbb{F}) \cup \mathrm{bn}(\mathbb{G})$ and $\mathrm{extr}(\mathbb{F}) = (\mathbb{F}_1, \mathbb{F}_2)$.

$$\mathbb{E}[\mathbb{F}[\overline{a}\langle Q\rangle] \parallel \mathbb{G}[a(X).P]] \to_{\mathrm{rs}} \mathbb{E}[\mathbb{F}_2[\mathbb{F}_1[\mathbf{0}] \parallel \mathbb{G}[P\{Q/X\}]]]$$
$$\mathbb{E}[\mathbb{G}[a(X).P] \parallel \mathbb{F}[\overline{a}\langle Q\rangle]] \to_{\mathrm{rs}} \mathbb{E}[\mathbb{F}_2[\mathbb{G}[P\{Q/X\}] \parallel \mathbb{F}_1[\mathbf{0}]]]$$

4.2 Zipper Semantics and NDAM

We present the zipper semantics of HOπ in Figure 5. The out and in transitions differ from HOcore as they carry two contexts $\mathbb{F}_1$ and $\mathbb{F}_2$: as in the reduction semantics, $\mathbb{F}_1$ collects the parallel compositions (rules outParL and outParR) while $\mathbb{F}_2$ collects the name restrictions (rule outNu).

Checking that the name a on which the communication happens is not captured by $\mathbb{F}_2$ or $\mathbb{G}$ is not done the same way in the out and in transitions, because the transitions themselves are not completely symmetric. In the input transition, we already know the name a, so we simply verify that the names bound by $\mathbb{G}$ differ from a on the fly in rule inNu. We cannot do the same in rule outNu, because we do yet not know a at this point. We know a when we find the output (rule outIn), so we check here that $\mathbb{F}_2$ does not capture it.

The zipper and reduction semantics coincide, the difficulty in the proof is in handling scope extrusion (cf. Appendix C).

Theorem 4.1. *For all P and P', $P \to_{\mathrm{rs}} P'$ iff $P \to_{\mathrm{zs}} P'$.*

Remark 3. The zipper semantics for HOπ cannot be written in the left-first style (Remark 2) because of scope extrusion. After finding the communicating processes $P \parallel Q$, we search for an output or input in P. Because we do not know the operator in advance, we do not know if we should decompose the context surrounding it to account for scope extrusion.

While writing the zipper semantics for HOπ requires some care, the corresponding NDAM is as expected. A difference with HOcore is the side-conditions in the outIn and inNu rules, which are added to the step, e.g., for outIn:

$$\langle \overline{a}^{\Sigma}\langle P\rangle ; \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R\rangle_{\mathsf{out}} \mapsto \langle R ; (\mathsf{outIn}, \Sigma) :: \pi \mid \bullet, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2\rangle_{\mathsf{in}} \quad \text{if } (\mathsf{in}, a) \notin \mathrm{an}(R) \text{ and } a \notin \mathrm{bn}(\mathbb{F}_2)$$

If the side-condition is not met, the "otherwise" step applies and we switch to the backward mode bout. The side-condition also makes the output mode annotation become $(\mathsf{out}, |R|, \mathbb{F}_2)$: a process $\overline{a}\langle P\rangle$ is a normal form w.r.t. output if $\mathbb{F}_2$ captures a, so being a normal form in this mode depends on $\mathbb{F}_2$. The complete machine can be found in Appendix C.

5 Derivation of the Abstract Machine

We show how to derive an abstract machine from a zipper semantics under some conditions. To this end, we specify zipper semantics as a transition system [20], a framework used to describe rule formats.

5.1 Zipper Semantics as a Transition System

Given an entity e, we write $\widetilde{e}$ for a possibly empty sequence $(e_1, \ldots, e_n)$ for some n. We assume a set $\mathcal{S}$ of sorts ranged over by s, denoting what we need to express the semantics of the language (contexts, names, etc), and which includes the sort t of terms that are reduced. For each sort s, let $\mathcal{O}_s$ be the signature of s, i.e., a set of operators, each having a typing $\widetilde{s} \to s$. In particular, we let op range over the operators of the terms $\mathcal{O}_t$. We also assume a set $\mathcal{F}$ of auxiliary functions that are used to build terms, like term substitution or context plugging, each of type $\widetilde{s} \to t$.

For each s, we assume an infinite set $\mathcal{V}_s$ of *rule variables*, denoted by v_s, w_s, or v, w if the sort does not matter. The set $\mathfrak{E}_s$ of *rule entities* of sort s, ranged over by e_s, f_s (or e, f if we ignore the sort), are the entities built out of the signature $\mathcal{O}_s$ extended with rule variables. We define $\mathfrak{E}_s$ inductively so that $\mathcal{V}_s \subseteq \mathfrak{E}_s$, and for all $o \in \mathcal{O}_s$ of signature $(s_1, \ldots, s_n) \to s$ and $(e_{s_i} \in \mathfrak{E}_{s_i})_{i \in 1\ldots n}$ for some n, we have $o(e_{s_1}, \ldots, e_{s_n}) \in \mathfrak{E}_s$. A special case are term entities e_t, which can also be built out of auxiliary functions in $\mathcal{F}$. We write $\mathrm{rv}(e_s)$ for the set of rule variables of e_s; e_s is ground if $\mathrm{rv}(e_s) = \emptyset$.

A rule substitution σ is a sort-respecting mapping from rule variables to rule entities. It should not be confused with the substitution $\cdot\{\cdot/\cdot\}$ which may exist for terms and is considered an auxiliary function in $\mathcal{F}$. We write $v\sigma$ for the application of σ to v, and $e\sigma$—for its extension to rule entities, defined in the expected way. A ground entity e is an instance of e' if there exists σ such that $e'\sigma = e$.

Given some rule variables $\widetilde{v}$, we write $\mathcal{P}(\widetilde{v})$ for a decidable predicate on $\widetilde{v}$. We assume a set $\mathcal{M}$ of modes, denoted by m, such that each mode is associated with a sequence $\widetilde{s_{\mathsf{m}}}$, which

$$\text{toy}\ \frac{P \xrightarrow{a,\mathbb{E}} P'}{P \xrightarrow{\mathbb{E}} P'} \qquad \text{outInL}\ \frac{R \xrightarrow{\mathbb{F}[\mathbf{0}] \parallel :: \mathbb{E},a,P}_{\mathsf{in}} P'}{\overline{a}\langle P\rangle \xrightarrow{\mathbb{F},\mathcal{L},\mathbb{E},R}_{\mathsf{out}} P'} \qquad \text{choiceBad}\ \frac{P \xrightarrow{\mathbb{E}} P'}{P + Q \xrightarrow{\mathbb{E}} P'} \qquad \text{choiceOk}\ \frac{P \xrightarrow{\mathbb{E},Q :: \theta} P'}{P + Q \xrightarrow{\mathbb{E},\theta} P'} \qquad \text{rec}\ \frac{P\{\mu X.P/X\} \xrightarrow{\mathbb{E}} P'}{\mu X.P \xrightarrow{\mathbb{E}} P'}$$

Figure 6. Rules for variants of HOcore

are the sorts of the arguments of m. The set $\mathcal{M}$ includes the initial mode zs associated to the empty sequence.

A *transition* is a predicate $e_i \xrightarrow{\widetilde{e}}_{\mathsf{m}} e_o$, where e_i and e_o are respectively the source and the target. We consider only three kinds of rule: inductive (whose names are ranged over with ρ), axiom, and initial, of the following respective shapes.

$$\frac{e_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}\ \rho \qquad \frac{\mathcal{P}(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} e_t} \qquad \frac{v_t \xrightarrow{\widetilde{f}}_{\mathsf{m}} w_t}{v_t \rightarrow_{\mathsf{zs}} w_t}\ \mathsf{init}$$

We extend the notion of set of rule variables rv and the application of a substitution to transitions and rules.

An inductive rule has only one premise, and may have side-conditions, represented by $\mathcal{P}$, on some of the variables $\widetilde{w}$ occurring in the rule. The modes m and m′ may be distinct or not, and the sequences $\widetilde{e}$ and $\widetilde{f}$ should be rule entities of sorts respectively $\widetilde{s_{\mathsf{m}}}$ and $\widetilde{s_{\mathsf{m}'}}$. The sources and targets of the transitions are terms; in the conclusion, the source term is of the form $op(\widetilde{v})$, enforcing that a rule can only pattern-match the head operator of the term. Both targets should be the same term variable, meaning that an inductive rule is simply passing along the result. Computation occurs in axioms, where the target can be any term.

An initial rule defines the initial mode zs. The source of the conclusion is a variable, so an initial rule does not perform any pattern-matching. An initial rule is just a means to set up the arguments of another mode m (such that $\mathsf{m} \neq \mathsf{zs}$). A zipper semantics is a triple $(\mathcal{S}, \mathcal{O}, \mathcal{R})$ where $\mathcal{R}$ is a finite set of zipper rules with exactly one initial rule. The associated semantics on terms is defined by $\rightarrow_{\mathsf{zs}}$.

5.2 Derivable Zipper Semantics

Not every zipper semantics can be turned into an NDAM. Some conditions have to be satisfied for the transformation to be possible and to ensure termination.

The first one is that the rules of the semantics must be constructive w.r.t. the machine, meaning that the entities in its premise are constructed from the ones in the conclusion. Indeed, the abstract machine searches for redexes with forward steps by going from the conclusion to the premise of a rule. As a result, a rule like toy in Figure 6 cannot be turned into a machine step, as the machine would have to guess the name a. We forbid such a rule by requiring that in each inductive rule of the zipper semantics, the rule variables of the premise are included in the rule variables of the conclusion.

Definition 5.1. $\dfrac{e_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$ is machine constructive if $\mathsf{rv}(e_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'}) \cup \widetilde{w} \subseteq \mathsf{rv}(op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}})$.

The other constraint is that the rules must be *reversible* to allow for backtracking: it should be possible to reconstruct the entities in the conclusion from the ones in the premise. We say a rule is reversible if it cannot have two different instances with the same premise. For example, we could make the input search in HOcore less verbose, by combining the contexts $\mathbb{E}$ and $\mathbb{F}$ in a single context, like in the rule outInL in Figure 6. In $\mathbb{E}[R \parallel \mathbb{F}[\mathbf{0}]]$, the input process is plugged into the context $\parallel \mathbb{F}[\mathbf{0}] :: \mathbb{E}$, that we build in rule outInL, instead of keeping $\mathbb{E}$ and $\mathbb{F}$ separate as in Figure 3. However, to unapply the rule outInL, we need to uniquely decompose a context as $\parallel \mathbb{F}[\mathbf{0}] :: \mathbb{E}$, which is not possible as soon as there are several occurrences of $\mathbf{0}$ in $\mathbb{F}[\mathbf{0}]$: the rule outInL is not reversible. We give a simple sufficient criterion for a rule to be reversible.

Lemma 5.2. $\dfrac{e_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$ *is reversible if we have* $\mathsf{rv}(op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}}) \subseteq \mathsf{rv}(e_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'})$, *and the auxiliary functions used to build the entities in* e_t *and* $\widetilde{f}$ *are injective.*

The first condition states that the rules variables of the conclusion have to be included in those of the premise. Indeed, if we forget an entity between the conclusion and the premise, like Q in the rule for choice choiceBad in Figure 6, then we have no information to restore Q when backtracking. Instead, it should be kept in an extra argument of the zipper semantics, like the stack θ in the rule choiceOk in Figure 6. The stack θ is useful only for backtracking and not to define the semantics of the language, as it is simply thrown away when we apply an axiom. Any rule forgetting entities between its conclusion and premise can be made reversible using this principle [37].

Finally, we want the machine to always terminate when searching for a redex. Consider for instance the rec rule for a recursion operator in Figure 6. The corresponding machine would infinitely loop with $\mu X.X$. Indeed, the forward step of this rule changes focus from the source of the conclusion to the source of the premise, but these two terms are equal when $P = X$. To avoid this, we require the zipper semantics to be well-founded in the following sense.

Definition 5.3. A zipper semantics is well-founded if there exists a well-founded size ζ on transitions such that for all inductive rules $\dfrac{e_t' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$, we have $\zeta(e_t' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t) < \zeta(e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t)$.

In the calculi of this paper, each rule either focuses on a subterm or it changes mode (like in rule outIn in HOcore). We therefore define an ordering on modes such that $\mathsf{m} > \mathsf{m}'$ if the derivation of m depends on m'; e.g., we have $\mathsf{zs} > \mathsf{app} > \mathsf{lam}$ in λ-calculus, and $\mathsf{zs} > \mathsf{par} > \mathsf{out} > \mathsf{in}$ in HOcore and HOπ. The size we consider is then the lexicographic ordering composed of the ordering on modes followed by the subterm ordering on the source term of the transition. This size works as long as we have no cyclic dependencies in modes and only congruence rules within each mode. It rules out unconstrained recursion, but we can still adapt it for guarded recursion, where the recursion variable occurs only after an input, as in $\mu X.a(Y).(X \parallel Y)$. In the premise of the rec rule, the μ operator itself becomes guarded, so the number of recursion operators at toplevel strictly decreases.

The semantics of Figures 1, 3, and 5 are machine constructive well-founded, and reversible (they satisfy Lemma 5.2). Henceforth, we assume the zipper semantics to be machine constructive, reversible, and well-founded.

5.3 Machine Derivation

Annotations. The machine annotates terms which cannot do certain transitions, to forbid repeated tries which would lead to an infinite loop. The arguments of the transition may play a role in whether the term is a normal form or not: in HOcore an output $\overline{a}\langle P\rangle$ is a normal form w.r.t. the output transition $\xrightarrow{\mathbb{F},\mathbb{S},\mathbb{E},R}_{\mathsf{out}}$ if R cannot receive the message on a, so the annotation is $(\mathsf{out}, |R|)$. Similarly an input $\xrightarrow{\mathbb{G},\mathbb{S},a,\mathbb{E},\mathbb{F}}_{\mathsf{in}}$ depends on the name a.

The arguments kept in the annotation are the ones either taking part in the reduction, like R in the output case, or in side-conditions, like a in the input case. Given a mode m with arguments $\widetilde{e}$, its annotation $\phi(\mathsf{m}, \widetilde{e})$ is defined as $(\mathsf{m}, \widetilde{f})$ where $\widetilde{f} \subseteq \widetilde{e}$ are the arguments occurring either in side-conditions or source terms of the rules defining m. Repeating this for each mode of a zipper semantics, we define the *annotation function* ϕ of the semantics.

Annotated terms. Let $(\mathcal{S}, \mathcal{O}, \mathcal{R})$ be a zipper semantics with annotation function ϕ. We extend $\mathcal{S}$ with the sort of annotation sets s_Σ, for which we assume the usual operators on sets. The machine is built on a signature $\mathcal{A}$ which replaces the signature for terms $\mathcal{O}_t$ with *annotated terms*, so that for all $op \in \mathcal{O}_t$ of type $(s_1, \ldots, s_n) \to t$ for some n, we have a corresponding operator $op \in \mathcal{A}_t$ of type $(s_\Sigma, s_1, \ldots, s_n) \to t$.

We let a range over annotated terms $\mathfrak{A}_t$, built out of $\mathcal{A}$, $\mathcal{V}_t$, and a single rule variable for annotation sets v_Σ: one variable is enough, as at most one annotation set occurs in a given machine step. Given an annotated term $a = op(e_\Sigma, \widetilde{e})$, we write $\mathsf{an}(a)$ for its annotation set e_Σ. Given a term $e_t \in \mathfrak{E}_t$, its annotated version, written $\|e_t\|$, is inductively defined so that $\|v_s\| = v_s$ and $\|op(\widetilde{e})\| = op(v_\Sigma, \widetilde{\|e\|})$. Given an annotated term $a \in \mathfrak{A}_t$, its erasure $|a|$ produces a term with empty annotation sets, inductively defined so that $|v_s| = v_s$ and $|op(e_\Sigma, \widetilde{e})| = op(\emptyset, \widetilde{|e|})$.

Machine steps. The syntax of rule stacks π is given by $\pi ::= \mathsf{init} \mid (\rho, \Sigma) :: \pi$. We denote configurations $\langle a \,;\, \pi \mid \widetilde{e}\rangle_{\mathsf{m}}$ as *forward*, a special case being *initial* ones $\langle a\rangle_{\mathsf{zs}}$. *Backward* configurations are of the form $\langle \pi \,;\, a \mid \widetilde{e}\rangle_{\mathsf{bm}}$ with *normal-form* ones $\langle a\rangle_{\mathsf{nf}}$ as a subcase.

An inductive rule and an initial rule

$$\frac{e_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}\ \rho \qquad\qquad \frac{v_t \xrightarrow{\widetilde{f}}_{\mathsf{m}} w_t}{v_t \longrightarrow_{\mathsf{zs}} w_t}\ \mathsf{init}$$

generate respectively the steps

$$\begin{aligned}
\langle op(v_\Sigma, \widetilde{v}) \,;\, \pi \mid \widetilde{e}\rangle_{\mathsf{m}} &\mapsto \langle \|e_t\| \,;\, (\rho, v_\Sigma) :: \pi \mid \widetilde{f}\rangle_{\mathsf{m}'} \\
&\qquad \text{if } \phi(\mathsf{m}', \widetilde{f}) \notin \mathsf{an}(\|e_t\|) \text{ and } \mathcal{P}(\widetilde{w}) \\
\langle (\rho, v_\Sigma) :: \pi \,;\, \|e_t\| \mid \widetilde{f}\rangle_{\mathsf{bm}'} &\mapsto \langle op(v_\Sigma, \widetilde{v}) \,;\, \pi \mid \widetilde{e}\rangle_{\mathsf{m}} \\
\langle v_t\rangle_{\mathsf{zs}} &\mapsto \langle v_t \,;\, \mathsf{init} \mid \widetilde{f}\rangle_{\mathsf{m}} \\
\langle \mathsf{init} \,;\, v_t \mid \widetilde{f}\rangle_{\mathsf{bm}} &\mapsto \langle v_t\rangle_{\mathsf{nf}}
\end{aligned}$$

The forward step for an inductive rule goes from the conclusion to the premise, while the backward step goes in the opposite direction. Terms are extended with the rule variable for annotated sets v_Σ. The initial rule case is the same as the inductive one but simpler, as there is no side-condition: the annotated sets of the term v_t are assumed to be empty.

An axiom $\dfrac{\mathcal{P}(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} e_t}$ generates the forward step

$$\langle op(v_\Sigma, \widetilde{v}) \,;\, \pi \mid \widetilde{e}\rangle_{\mathsf{m}} \mapsto \langle \mid \|e_t\| \mid \rangle_{\mathsf{zs}} \text{ if } \mathcal{P}(\widetilde{w})$$

The annotations are erased after applying an axiom. There is no backward step associated with axioms.

What remains are the switching steps from a forward to a backward configuration, i.e., when we realize that the current mode m does not apply to the term $op(v_\Sigma, \widetilde{v})$ we reduce. These are the "otherwise" steps in Figures 2 and 4, which actually cover different cases. The first possibility is that op does not have a rule applying to it in the mode m. For such cases, we add a step

$$\langle op(v_\Sigma, \widetilde{v}) \,;\, \pi \mid \widetilde{e}\rangle_{\mathsf{m}} \mapsto \langle \pi \,;\, op(v_\Sigma \cup \{\phi(\mathsf{m}, \widetilde{e})\}, \widetilde{v}) \mid \widetilde{e}\rangle_{\mathsf{bm}}$$

When going to a backward configuration, we extend the annotation set of the operator with the current annotation.

The other case is that no rule $\dfrac{e_t^i \xrightarrow{\widetilde{f_i}}_{\mathsf{m}_i} v_t \qquad \mathcal{P}_i(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}\ \rho_i$ for op in the mode m applies, because either the premise or the side condition do not hold. If the machine has already checked that the premise fails, then e_t^i has been annotated with $\phi(\mathsf{m}_i, \widetilde{f_i})$. The corresponding switching step is therefore

$$\langle op(v_\Sigma, \widetilde{v}) ; \pi \mid \widetilde{e} \rangle_{\mathsf{m}} \mapsto \langle \pi ; op(v_\Sigma \cup \{\phi(\mathsf{m}, \widetilde{e})\}, \widetilde{v}) \mid \widetilde{e} \rangle_{\mathsf{bm}} \quad \text{if} \bigwedge_i \Big(\phi(\mathsf{m}_i, \widetilde{f_i}) \in \mathsf{an}(\|e_t^i\|) \vee \neg \mathcal{P}_i(\widetilde{w})\Big)$$

6 Correspondence Results

We state the correspondence between a machine constructive, reversible, well-founded zipper semantics and its derived machine. We prove that the search for a redex terminates (Theorem 6.3), and that the result is either an initial or a normal-form configuration (Theorem 6.4).

These outcomes correspond to respectively a zipper transition or lack thereof (Theorem 6.9). The difficult part is to show that a machine run corresponds to a zipper derivation. The intuition is that a machine run is well-bracketed, in the sense that backward steps undo the rules applied by forward steps in the reverse order: we can witness this behavior in Example 2. If the machine ends up applying an axiom, then we can read from the machine run the derivation tree of the corresponding transition in zipper semantics. The proofs of this section and auxiliary lemmas can be found in Appendix D.

6.1 Preliminary Notations

To distinguish between rule entities e_t—used to abstract semantics rules and machine steps—and their instances $e_t\sigma$, we write the latter using capital letters. In particular, we write T for ground terms and A for annotated ground terms. Given a ground term T, we write $\|T\|^\emptyset$ for the ground annotated term with empty annotation sets, i.e., $\|T\|^\emptyset = \|T\|\{v_\Sigma \mapsto \emptyset\}$. For all A, there exists a unique T such that $|A| = \|T\|^\emptyset$. In statements relating the machine to the zipper semantics, we identify $|A|$ and T, writing, e.g., $\vdash |A| \xrightarrow{\widetilde{E}}_{\mathsf{m}} T$.

We let $\mathcal{C}$, $\mathcal{F}$, $\mathcal{I}$, and $\mathcal{B}$ range over respectively all, forward, initial, and backward configurations. We write $\mapsto$ for the stepping relation between configurations, and $\mapsto^*$ and $\mapsto^+$ for its (reflexive and) transitive closures, respectively. A *search path* is a sequence $\mathcal{C} \mapsto^* \mathcal{C}'$ where only $\mathcal{C}$ or $\mathcal{C}'$ may be initial configurations.

We sometimes make explicit the one-to-one correspondence between rules and machine steps by labeling a forward step $\mathcal{F} \overset{\rho}{\mapsto} \mathcal{F}'$ and backward step $\mathcal{B} \overset{-\rho}{\longmapsto} \mathcal{F}$ with the rule ρ it applies or unapplies. We label a switching step τ ($\mathcal{F} \overset{\tau}{\mapsto} \mathcal{B}$), and let $\lambda ::= \rho \mid -\rho \mid \tau$.

An arbitrary configuration may contain annotations inconsistent with a machine run (like $\lambda^{\mathsf{lam}} x.t$). To rule out such configurations, we define validity as follows.

Definition 6.1. A configuration $\mathcal{C}$ is valid if there exists a ground term T such that $\langle \|T\|^\emptyset \rangle_{\mathsf{zs}} \mapsto^* \mathcal{C}$.

6.2 Properties of the NDAM

Before stating the correspondence results themselves, we prove some properties of the NDAM, like its termination.

We first prove that the annotations behave as expected: if an annotated term ends up in a backward configuration, it is because it cannot do the corresponding zipper transition.

Lemma 6.2. *Let $\mathcal{B} = \langle \pi ; A \mid \widetilde{E} \rangle_{\mathsf{bm}}$ be a valid configuration. We have $\neg(|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}})$.*

The proof relies on the fact that annotations contains enough information to ensure that a term cannot do a transition: (cf. Lemma D.1 in the appendix).

Next, we show that the machine cannot loop in the search for a redex, i.e., search paths are finite. Roughly, we cannot have an infinite sequence of forward steps $\overset{\rho}{\mapsto}$, because the well-founded hypothesis on the zipper semantics carries over to these steps. We cannot have an infinite sequence of switching $\overset{\tau}{\mapsto}$ steps, because each of them adds one annotation, and the number of annotations is bounded for a given term (see Lemma D.2 in the appendix). By construction, backward steps $\overset{-\rho}{\longmapsto}$ can only follow a switching step.

Theorem 6.3. *For all $\mathcal{I}$, there exists a number n such that any search path starting from $\mathcal{I}$ has length at most n.*

Finally, the possible outcomes of a search path are either in an initial or a normal-form configuration.

Theorem 6.4. *Let $\mathcal{C}$ be a valid configuration. Either $\mathcal{C} \mapsto^* \mathcal{I}$ for some $\mathcal{I}$, or $\mathcal{C} \mapsto^* \langle A \rangle_{\mathsf{nf}}$ for some A.*

6.3 Correspondence Results

In what follows, we write $\vdash T \to_{\mathsf{zs}} T'$ if there exists a zipper derivation with $T \to_{\mathsf{zs}} T'$ as conclusion; see Appendix D for the formal definition.

A zipper transition can be reflected as a run which never backtracks, just by mimicking the derivation tree.

Theorem 6.5. *For all $\vdash T \to_{\mathsf{zs}} T'$, we have $\langle \|T\|^\emptyset \rangle_{\mathsf{zs}} \mapsto^* \langle \|T'\|^\emptyset \rangle_{\mathsf{zs}}$.*

We show that the machine is also complete w.r.t. normal forms, i.e., if a term cannot reduce in the semantics, the machine ends in the machine state for normal forms $\langle \cdot \rangle_{\mathsf{nf}}$.

Theorem 6.6. *If $\neg(T \to_{\mathsf{zs}})$, then any machine run from the configuration $\langle \|T\|^\emptyset \rangle_{\mathsf{zs}}$ ends with $\langle A \rangle_{\mathsf{nf}}$ such that $|A| = T$.*

The machine goes through a normal form to annotate each of its constructors from which a transition could trigger. Different runs produce the same annotations for each constructor,

but generated in a different order, depending on the non-deterministic choices the machine makes.

To prove that a search path encodes a zipper derivation, we formalize the idea that a machine run is well-bracketed: backward steps undo the rules in the reverse order. We define the backtrack-free closure $\overset{\rho}{\mapsto}_{\mathsf{bf}}$ of a search path, which forgets about the rules that have been applied and then unapplied. We write $\overset{\lambda_1.\lambda_2 \ldots \lambda_n}{\longmapsto}_{\mathsf{bf}}$ for a sequence $\overset{\lambda_1}{\longmapsto}_{\mathsf{bf}} \overset{\lambda_2}{\longmapsto}_{\mathsf{bf}} \ldots \overset{\lambda_n}{\longmapsto}_{\mathsf{bf}}$. The $\mapsto_{\mathsf{bf}}$ relation is defined as follows:

$$\frac{\mathcal{C} \overset{\lambda}{\mapsto} \mathcal{C}'}{\mathcal{C} \overset{\lambda}{\mapsto}_{\mathsf{bf}} \mathcal{C}'} \qquad \frac{\mathcal{F} \overset{\rho.\tau.-\rho}{\longmapsto}_{\mathsf{bf}} \mathcal{F}'}{\mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{F}'} \qquad \frac{\mathcal{F} \overset{\tau.\tau}{\longmapsto}_{\mathsf{bf}} \mathcal{F}'}{\mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{F}'} \qquad \frac{\mathcal{F} \overset{\tau.\rho}{\longmapsto}_{\mathsf{bf}} \mathcal{F}'}{\mathcal{F} \overset{\rho}{\mapsto}_{\mathsf{bf}} \mathcal{F}'}$$

Backward-free steps extend regular steps with a new behavior, as we can have a $\overset{\tau}{\mapsto}_{\mathsf{bf}}$ step between forward configurations. In such a case, the two configurations are equal up to annotations: the resulting configuration contains strictly more annotations than the source one. The other backtrack-free steps corresponds to regular machine steps.

A run can be represented as a sequence of backtrack-free steps (Lemma D.8), and if such sequence ends with an axiom application, it mimics a derivation tree.

Theorem 6.7. *For all search path* $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}} \mapsto^{+} \langle A \rangle_{\mathsf{zs}}$*, we have* $\vdash T \to_{\mathsf{zs}} |A|$.

The machine is also sound w.r.t. normal forms: if a run ends with $\langle A \rangle_{\mathsf{nf}}$, then $|A|$ is a normal form. This property is a direct consequence of Lemma 6.2, by taking the initial mode.

Theorem 6.8. *Let* $\langle A \rangle_{\mathsf{nf}}$ *be a valid configuration; then* $|A|$ *is a normal form.*

Combining the results together, the correspondence between a zipper semantics and its NDAM is as follows.

Theorem 6.9. *For all* T, T', *and* A,

- $\vdash T \to_{\mathsf{zs}} T'$ *iff there exists a search path* $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}} \mapsto^{+} \langle \|T'\|^{\emptyset} \rangle_{\mathsf{zs}}$;
- *the term* T *is a normal form iff there exists a search path* $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}} \mapsto^{+} \langle A \rangle_{\mathsf{nf}}$ *with* $|A| = T$.

7 Related Work

The zipper semantics of the process calculi are inspired by complementary semantics [30], a format dedicated to bisimulation-based equivalence proofs. In both semantics, the derivation tree of two communicating processes is sequentialized. The difference is in the transition labels, which should be as minimal as possible in complementary semantics to keep the bisimulation proofs simple, while ours are detailed enough to be able to reconstruct the whole term.

Typical abstract machines for deterministic languages based on the λ-calculus are in refocused form [9]; such machines continue term decomposition from the contraction site. They have been shown to be uniformly derivable from the underlying reduction semantics by a refocusing method [5, 42], and the correctness of the derivation hinges on the unique decomposition property. In contrast, NDAMs do not have this property, and after contracting a redex they completely reconstruct the term to start decomposition again. An optimization similar to refocusing for non-deterministic languages appears more challenging in general. Another common feature of abstract machines for the λ-calculus is an efficient implementation of substitution with environments [6]. The use of environments is orthogonal to the derivation of NDAMs: if the source zipper semantics uses environments, then so does its derived NDAM. We consider substitution-based zipper semantics in this paper because they are simpler than environments-based ones.

Process algebras have been implemented in various frameworks ranging from rewriting logic [43] to biological systems [33], including dedicated implementations and abstract machines [4, 14, 16–19, 22, 32, 35, 36, 44]. These implementations are ad-hoc and calculus-specific, and only some of them are complete [4, 17, 19, 22, 35, 36]. We believe we can handle most of these calculi in our framework in a uniform and complete way. However, the resulting implementation would be "single-threaded", while the distribution of processes is a concern of previous machines [17], especially for calculi with localities [4, 18, 19, 22, 36]. Considering the many different models of distribution, making our machine distributed requires significantly more work, especially if we want to remain generic and complete.

Our use of backtracking evokes reversible calculi [7, 45], where one can revert communication steps, not necessarily in the order they happened, as long as the causality between them is preserved. The concerns are different, however: theirs is to keep enough information to track causality [28, 37], while ours is to control backtracking to avoid infinite searches. As a result, we store less information in machine configurations, but the annotations we use to prevent loops are typically not needed in their setting.

8 Conclusion

We present a generic design of abstract machines for non-deterministic languages. The machine navigates a term in the search for a redex, making arbitrary choices when several paths are possible, and backtracks when it reaches a subterm which cannot reduce. The machine annotates such subterms and thus avoids trying again an already explored path, which prevents infinite search. An NDAM is automatically derived from zipper semantics, a form of SOS in which the decomposition process of a term into a context and a redex is made explicit. The machine is sound and complete w.r.t. the zipper semantics. The derivation procedure has been implemented in OCaml [1]. The presented methodology is readily applicable to other non-deterministic calculi

not shown in this paper, such as concurrent lambda calculi, with communication via channels or via futures [3, 15, 34].

An improvement of the current design would be to keep as many annotations as possible when reducing, rather than erasing the whole term, in order to prune redundant search. Another optimization would be to find a way to manage annotations that would generically enable refocusing.

In terms of expressiveness, we think that a different design of zipper semantics, where we search for constructs at the leaves of the redex syntax tree instead of starting at the root, could better handle calculi which cannot be easily expressed with a reduction semantics, like the Join calculus [14, 16].

A final question is whether the zipper semantics could be itself derived from a more commonly used format, such as context-based reduction semantics or SOS. Finding the appropriate starting point is already an issue in itself, as it should be able to express the different families of non-deterministic languages, such as concurrent λ-calculi or process calculi.

References

[1] Non-deterministic abstract machines. Implementation available at https://link.infini.fr/ndamlics22.

[2] M. S. Ager, D. Biernacki, O. Danvy, and J. Midtgaard. A functional correspondence between evaluators and abstract machines. In *Proceedings of the 5th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, 27-29 August 2003, Uppsala, Sweden*, pages 8–19. ACM, 2003.

[3] F. Aschieri, A. Ciabattoni, and F. A. Genco. On the concurrent computational content of intermediate logics. *Theor. Comput. Sci.*, 813:375–409, 2020.

[4] P. Bidinger, A. Schmitt, and J. Stefani. An abstract machine for the kell calculus. In M. Steffen and G. Zavattaro, editors, *Formal Methods for Open Object-Based Distributed Systems, 7th IFIP WG 6.1 International Conference, FMOODS 2005, Athens, Greece, June 15-17, 2005, Proceedings*, volume 3535 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2005.

[5] M. Biernacka, W. Charatonik, and K. Zielinska. Generalized refocusing: From hybrid strategies to abstract machines. In D. Miller, editor, *2nd International Conference on Formal Structures for Computation and Deduction, FSCD 2017, September 3-9, 2017, Oxford, UK*, volume 84 of *LIPIcs*, pages 10:1–10:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[6] M. Biernacka and O. Danvy. A concrete framework for environment machines. *ACM Trans. Comput. Log.*, 9(1):6, 2007.

[7] V. Danos and J. Krivine. Reversible communicating systems. In P. Gardner and N. Yoshida, editors, *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings*, volume 3170 of *Lecture Notes in Computer Science*, pages 292–307. Springer, 2004.

[8] O. Danvy. From reduction-based to reduction-free normalization. In P. W. M. Koopman, R. Plasmeijer, and S. D. Swierstra, editors, *Advanced Functional Programming, 6th International School, AFP 2008, Heijen, The Netherlands, May 2008, Revised Lectures*, volume 5832 of *Lecture Notes in Computer Science*, pages 66–164. Springer, 2008.

[9] O. Danvy and L. R. Nielsen. Syntactic theories in practice. *Electron. Notes Theor. Comput. Sci.*, 59(4):358–374, 2001.

[10] M. Felleisen, R. B. Findler, and M. Flatt. *Semantics Engineering with PLT Redex*. The MIT Press, 2009.

[11] M. Felleisen and D. P. Friedman. Control operators, the SECD-machine, and the λ-calculus. In M. Wirsing, editor, *Formal Description of Programming Concepts - III: Proceedings of the IFIP TC 2/WG 2.2 Working Conference on Formal Description of Programming Concepts - III, Ebberup, Denmark, 25-28 August 1986*, pages 193–222. North-Holland, 1987.

[12] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theor. Comput. Sci.*, 103(2):235–271, 1992.

[13] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theor. Comput. Sci.*, 103(2):235–271, 1992.

[14] F. L. Fessant. *JoCaml: conception et implémentation d'un langage à agents mobiles*. PhD thesis, École polytechnique, 2001.

[15] C. Flanagan and M. Felleisen. The semantics of future and an application. *J. Funct. Program.*, 9(1):1–31, 1999.

[16] C. Fournet and G. Gonthier. The reflexive CHAM and the join-calculus. In H. Boehm and G. L. S. Jr., editors, *Conference Record of POPL'96: The 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium, St. Petersburg Beach, Florida, USA, January 21-24, 1996*, pages 372–385. ACM Press, 1996.

[17] P. Gardner, C. Laneve, and L. Wischik. The fusion machine. In L. Brim, P. Jancar, M. Kretínský, and A. Kucera, editors, *CONCUR 2002 - Concurrency Theory, 13th International Conference, Brno, Czech Republic,*

August 20-23, 2002, Proceedings, volume 2421 of *Lecture Notes in Computer Science*, pages 418–433. Springer, 2002.

[18] F. Germain, M. Lacoste, and J. Stefani. An abstract machine for a higher-order distributed process calculus. *Electron. Notes Theor. Comput. Sci.*, 66(3):145–169, 2002.

[19] P. Giannini, D. Sangiorgi, and A. Valente. Safe ambients: Abstract machine and distributed implementation. *Sci. Comput. Program.*, 59(3):209–249, 2006.

[20] J. F. Groote and F. W. Vaandrager. Structured operational semantics and bisimulation as a congruence. *Inf. Comput.*, 100(2):202–260, 1992.

[21] J. Hannan and D. Miller. From operational semantics for abstract machines. *Math. Struct. Comput. Sci.*, 2(4):415–459, 1992.

[22] D. Hirschkoff, D. Pous, and D. Sangiorgi. An efficient abstract machine for safe ambients. *J. Log. Algebraic Methods Program.*, 71(2):114–149, 2007.

[23] G. P. Huet. The zipper. *J. Funct. Program.*, 7(5):549–554, 1997.

[24] S. L. P. Jones. Implementing lazy functional languages on stock hardware: The spineless tagless G-machine. *J. Funct. Program.*, 2(2):127–202, 1992.

[25] C. Kirchner, R. Kopetz, and P. Moreau. Anti-pattern matching. In R. D. Nicola, editor, *Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Held as Part of the Joint European Conferences on Theory and Practics of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings*, volume 4421 of *Lecture Notes in Computer Science*, pages 110–124. Springer, 2007.

[26] J.-L. Krivine. A call-by-name lambda-calculus machine. *Higher-Order and Symbolic Computation*, 20(3):199–207, 2007.

[27] P. J. Landin. The mechanical evaluation of expressions. *The Computer Journal*, 6(4):308–320, 1964.

[28] I. Lanese and D. Medic. A general approach to derive uncontrolled reversible semantics. In I. Konnov and L. Kovács, editors, *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, volume 171 of *LIPIcs*, pages 33:1–33:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[29] I. Lanese, J. A. Pérez, D. Sangiorgi, and A. Schmitt. On the expressiveness and decidability of higher-order process calculi. In *Proceedings of the Twenty-Third Annual IEEE Symposium on Logic in Computer Science, LICS 2008, 24-27 June 2008, Pittsburgh, PA, USA*, pages 145–155. IEEE Computer Society, 2008.

[30] S. Lenglet, A. Schmitt, and J. Stefani. Characterizing contextual equivalence in calculi with passivation. *Inf. Comput.*, 209(11):1390–1433, 2011.

[31] X. Leroy. The ZINC experiment: an economical implementation of the ML language. Technical report 117, INRIA, 1990.

[32] L. M. B. Lopes, F. M. A. Silva, and V. T. Vasconcelos. A virtual machine for a process calculus. In G. Nadathur, editor, *Principles and Practice of Declarative Programming, International Conference PPDP'99, Paris, France, September 29 - October 1, 1999, Proceedings*, volume 1702 of *Lecture Notes in Computer Science*, pages 244–260. Springer, 1999.

[33] U. Majumder and J. H. Reif. Design of a biomolecular device that executes process algebra. *Nat. Comput.*, 10(1):447–466, 2011.

[34] J. Niehren, J. Schwinghammer, and G. Smolka. A concurrent lambda calculus with futures. *Theor. Comput. Sci.*, 364(3):338–356, 2006.

[35] A. Phillips and L. Cardelli. A correct abstract machine for the stochastic pi-calculus. In *Concurrent Models in Molecular Biology*, 2004.

[36] A. Phillips, N. Yoshida, and S. Eisenbach. A distributed abstract machine for boxed ambient calculi. In D. A. Schmidt, editor, *Programming Languages and Systems, 13th European Symposium on Programming, ESOP 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 - April 2, 2004, Proceedings*, volume 2986 of *Lecture Notes in Computer Science*, pages 155–170. Springer, 2004.

[37] I. C. C. Phillips and I. Ulidowski. Reversing algebraic process calculi. *J. Log. Algebraic Methods Program.*, 73(1-2):70–96, 2007.

[38] G. D. Plotkin. A structural approach to operational semantics. Technical Report FN-19, DAIMI, Department of Computer Science, Aarhus University, Aarhus, Denmark, Sept. 1981.

[39] S. Salvati and I. Walukiewicz. Krivine machines and higher-order schemes. *Inf. Comput.*, 239:340–355, 2014.

[40] D. Sangiorgi. Bisimulation in higher-order process calculi. In E. Olderog, editor, *Programming Concepts, Methods and Calculi, Proceedings of the IFIP TC2/WG2.1/WG2.2/WG2.3 Working Conference on Programming Concepts, Methods and Calculi (PROCOMET '94) San Miniato, Italy, 6-10 June, 1994*, volume A-56 of *IFIP Transactions*, pages 207–224. North-Holland, 1994.

[41] D. Sangiorgi and D. Walker. *The Pi-Calculus - a theory of mobile processes*. Cambridge University Press, 2001.

[42] F. Sieczkowski, M. Biernacka, and D. Biernacki. Automating derivations of abstract machines from reduction semantics: - A generic formalization of refocusing in Coq. In J. Hage and M. T. Morazán, editors, *Implementation and Application of Functional Languages - 22nd International Symposium, IFL 2010, Alphen aan den Rijn, The Netherlands, September 1-3, 2010, Revised Selected Papers*, volume 6647 of *Lecture Notes in Computer Science*, pages 72–88. Springer, 2010.

[43] P. Thati, K. Sen, and N. Martí-Oliet. An executable specification of asynchronous pi-calculus semantics and may testing in maude 2.0. *Electron. Notes Theor. Comput. Sci.*, 71:261–281, 2002.

[44] D. Turner. *The Polymorphic Pi-calculus:Theory and Implementation*. PhD thesis, University of Edinburgh, 1995.

[45] I. Ulidowski, I. Lanese, U. P. Schultz, and C. Ferreira, editors. *Reversible Computation: Extending Horizons of Computing - Selected Results of the COST Action IC1405*, volume 12070 of *Lecture Notes in Computer Science*. Springer, 2020.

$$\text{init}\ \frac{P \xrightarrow{\bullet}_{\mathsf{par}} P'}{P \to_{\mathsf{zs}} P'} \qquad \text{parL}\ \frac{P \xrightarrow{\parallel Q :: \mathbb{E}}_{\mathsf{par}} P'}{P \parallel Q \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'}\ (s) \qquad \text{parCom}\ \frac{P \xrightarrow{\bullet,\mathbb{E},Q}_{\mathsf{left}} P'}{P \parallel Q \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'}$$

$$\text{leftParL}\ \frac{P \xrightarrow{\parallel Q :: \mathbb{F},\mathbb{E},R}_{\mathsf{left}} P'}{P \parallel Q \xrightarrow{\mathbb{F},\mathbb{E},R}_{\mathsf{left}} P'}\ (s) \qquad \text{leftOut}\ \frac{R \xrightarrow{\bullet,a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} P'}{\overline{a}\langle P\rangle \xrightarrow{\mathbb{F},\mathbb{E},R}_{\mathsf{left}} P'} \qquad \text{leftIn}\ \frac{R \xrightarrow{\bullet,a,X,P,\mathbb{E},\mathbb{F}}_{\mathsf{out}} P'}{a(X).P \xrightarrow{\mathbb{F},\mathbb{E},R}_{\mathsf{left}} P'}$$

$$\text{inParL}\ \frac{R \xrightarrow{\parallel Q :: \mathbb{G},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} P'}{R \parallel Q \xrightarrow{\mathbb{G},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} P'}\ (s) \qquad \text{inCom}\ \frac{}{a(X).R \xrightarrow{\mathbb{G},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} \mathbb{E}[\mathbb{F}[\mathbf{0}] \parallel \mathbb{G}[R\{P/X\}]]}$$

$$\text{outParL}\ \frac{R \xrightarrow{\parallel Q :: \mathbb{G},a,X,P,\mathbb{E},\mathbb{F}}_{\mathsf{out}} P'}{R \parallel Q \xrightarrow{\mathbb{G},a,X,P,\mathbb{E},\mathbb{F}}_{\mathsf{out}} P'}\ (s) \qquad \text{outCom}\ \frac{}{\overline{a}\langle R\rangle \xrightarrow{\mathbb{G},a,X,P,\mathbb{E},\mathbb{F}}_{\mathsf{out}} \mathbb{E}[\mathbb{F}[P\{R/X\}] \parallel \mathbb{G}[\mathbf{0}]]}$$

Figure 7. Left-first Zipper Semantics for HOcore

A Lambda-calculus

Lemma A.1. *For all* $t \xrightarrow{s,\mathbb{E}}_{\mathsf{lam}} t'$, *we have* $\mathbb{E}[t\ @\ s] \to_{\mathsf{rs}} t'$.
For all $t \xrightarrow{\mathbb{E}}_{\mathsf{app}} t'$, *we have* $\mathbb{E}[t] \to_{\mathsf{rs}} t'$.

Proof. The first item is by definition, and the second one is proved by induction on the derivation of $t \xrightarrow{\mathbb{E}}_{\mathsf{app}} t'$. The base case is app$\beta$, where we conclude using the first item.

For the recursive case, suppose we apply appL: we have $t\ @\ s \xrightarrow{\mathbb{E}}_{\mathsf{app}} t'$ because $t \xrightarrow{@\ s :: \mathbb{E}}_{\mathsf{app}} t'$. By induction, we have $(@\ s :: \mathbb{E})[t] \to_{\mathsf{rs}} t'$, i.e., $\mathbb{E}[t\ @\ s] \to_{\mathsf{rs}} t'$, as wished. The other cases are similar. □

Theorem A.2. *For all* $t \to_{\mathsf{zs}} t'$, *we have* $t \to_{\mathsf{rs}} t'$.

For completeness, we want to prove that $\mathbb{E}[(\lambda x.t'')\ @\ s] \to_{\mathsf{rs}} \mathbb{E}[t''\{s/x\}]$ implies $\mathbb{E}[(\lambda x.t'')\ @\ s] \to_{\mathsf{zs}} \mathbb{E}[t''\{s/x\}]$. First, we notice that $\lambda x.t'' \xrightarrow{s,\mathbb{E}}_{\mathsf{lam}} \mathbb{E}[t''\{s/x\}]$ holds by definition of $\xrightarrow{s,\mathbb{E}}_{\mathsf{lam}}$. With app$\beta$, we get $(\lambda x.t'')\ @\ s \xrightarrow{\mathbb{E}}_{\mathsf{app}} \mathbb{E}[t''\{s/x\}]$. To conclude, we use the following result.

Lemma A.3. *For all* $t \xrightarrow{\mathbb{E}}_{\mathsf{app}} t'$, *we have* $\mathbb{E}[t] \xrightarrow{\bullet}_{\mathsf{app}} t'$.

Proof. We proceed by induction on $\mathbb{E}$. There is nothing to prove for $\bullet$. If $\mathbb{E} = \lambda x :: \mathbb{E}'$, then $t \xrightarrow{\lambda x :: \mathbb{E}'}_{\mathsf{app}} t'$ implies $\lambda x.t \xrightarrow{\mathbb{E}'}_{\mathsf{app}} t'$ by appλ, from which we deduce $\mathbb{E}'[\lambda x.t] \xrightarrow{\bullet}_{\mathsf{app}} t'$ by the induction hypothesis , i.e., $\mathbb{E}[t] \xrightarrow{\bullet}_{\mathsf{app}} t'$, as wished. The proof is similar in the remaining cases. □

Theorem A.4. *For all* $t \to_{\mathsf{rs}} t'$, *we have* $t \to_{\mathsf{zs}} t'$.

B HOcore

The output-first zipper semantics is equivalent to reduction semantics in the following way.

Lemma B.1. *For all* $R \xrightarrow{\mathbb{G},a,\mathcal{S},P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} R'$, *there exists* R'' *such that either* $R' = \mathbb{E}[\mathbb{F}[\mathbf{0}] \parallel \mathbb{G}[R''\{P/X\}]]$ *if* $\mathcal{S} = \mathcal{L}$ *or* $R' = \mathbb{E}[\mathbb{G}[R''\{P/X\}] \parallel \mathbb{F}[\mathbf{0}]]$ *if* $\mathcal{S} = \mathcal{R}$.
For all $P \xrightarrow{\mathbb{F},\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$, *we have either* $\mathbb{E}[\mathbb{F}[P] \parallel R] \to_{\mathsf{rs}} P'$ *if* $\mathcal{S} = \mathcal{L}$ *or* $\mathbb{E}[R \parallel \mathbb{F}[P]] \to_{\mathsf{rs}} P'$ *if* $\mathcal{S} = \mathcal{R}$.
For all $P \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'$, *we have* $\mathbb{E}[P] \to_{\mathsf{rs}} P'$.

Each result is proved by induction on the zipper derivation.

Theorem B.2. *For all $P \to_{\mathsf{zs}} P'$, we have $P \to_{\mathsf{rs}} P'$.*

The reverse implication relies on the following results about contexts in zipper semantics.

Lemma B.3. *For all $R \xrightarrow{\mathbb{G},\mathcal{S},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} R'$, we have $\mathbb{G}[R] \xrightarrow{\bullet,\mathcal{S},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}} R'$.*
For all $P \xrightarrow{\mathbb{F},\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$, we have $\mathbb{F}[P] \xrightarrow{\bullet,\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$.
For all $P \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'$, we have $\mathbb{E}[P] \xrightarrow{\bullet}_{\mathsf{par}} P'$.

Suppose $R \to_{\mathsf{rs}} R'$ with $R = \mathbb{E}[\mathbb{F}[\overline{a}\langle Q\rangle] \parallel \mathbb{G}[a(X).P]]$; the proof is similar in the symmetric case. We have $a(X).P \xrightarrow{\mathbb{G},\mathcal{L},a,Q,\mathbb{E},\mathbb{F}}_{\mathsf{in}} R'$, and by the first item of Lemma B.3, we deduce $\mathbb{G}[a(X).P] \xrightarrow{\bullet,\mathcal{L},a,Q,\mathbb{E},\mathbb{F}}_{\mathsf{in}} R'$. We get $\overline{a}\langle Q\rangle \xrightarrow{\mathbb{F},\mathcal{L},\mathbb{E},\mathbb{G}[a(X).P]}_{\mathsf{out}} R'$ by rule outIn, i.e., $\mathbb{F}[\overline{a}\langle Q\rangle] \xrightarrow{\bullet,\mathcal{L},\mathbb{E},\mathbb{G}[a(X).P]}_{\mathsf{out}} R'$ with the second item. With rule parOutL, we obtain $\mathbb{F}[\overline{a}\langle Q\rangle] \parallel \mathbb{G}[a(X).P] \xrightarrow{\mathbb{E}}_{\mathsf{par}} R'$, from which we can conclude using the last item.

Theorem B.4. *For all $P \to_{\mathsf{rs}} P'$, we have $P \to_{\mathsf{zs}} P'$.*

The left-first semantics for HOcore is given in Figure 7. The par transition is going through the process to find the parallel composition at the root of the communication redex, building the context $\mathbb{E}$ surrounding the redex at the same time. Finding the parallel composition triggers the $\xrightarrow{\mathbb{F},\mathbb{E},R}_{\mathsf{left}}$ transition, which looks for an input or an output in the process on the left, while building the context $\mathbb{F}$ and remembering $\mathbb{E}$ and the process on the right R. If we find an output, we look for an input on the same name in R using $\xrightarrow{\mathbb{G},a,P,\mathbb{E},\mathbb{F}}_{\mathsf{in}}$ (rule leftOut), otherwise we look for an output using $\xrightarrow{\mathbb{G},a,X,P,\mathbb{E},\mathbb{F}}_{\mathsf{out}}$ (rule leftIn). These two transitions are building the context $\mathbb{G}$ and use the remaining arguments to compute the results of the communication (rules inCom and outCom).

The corresponding NDAM is in Figure 8, except for the out and bout modes, which are symmetric to the in and bin modes.

C Correspondence results and NDAM for HOπ

We first prove that zipper semantics implies reduction semantics.

Lemma C.1. *For all transitions $R \xrightarrow{\mathbb{G},a,\mathcal{S},P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}} R'$, there exists R'' such that $R' = \mathbb{E}[\mathbb{F}_2[\mathbb{F}_1[\mathbf{0}] \parallel \mathbb{G}[R''\{P/X\}]]]$ if $\mathcal{S} = \mathcal{L}$ and $R' = \mathbb{E}[\mathbb{F}_2[\mathbb{G}[R''\{P/X\}] \parallel \mathbb{F}_1[\mathbf{0}]]]$ if $\mathcal{S} = \mathcal{R}$.*
For all transitions $P \xrightarrow{\mathbb{F}_1,\mathbb{F}_2,\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$ and $\mathbb{F}$ such that $\mathsf{extr}(\mathbb{F}) = (\mathbb{F}_1, \mathbb{F}_2)$, we have $\mathbb{E}[\mathbb{F}[P] \parallel R] \to_{\mathsf{rs}} P'$ if $\mathcal{S} = \mathcal{L}$ and $\mathbb{E}[R \parallel \mathbb{F}[P]] \to_{\mathsf{rs}} P'$ if $\mathcal{S} = \mathcal{R}$.
For all $P \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'$, we have $\mathbb{E}[P] \to_{\mathsf{rs}} P'$.

Proof. We sketch the proof of the second item, the others are easy. The proof is by induction on the derivation of the out transition. We assume $\mathcal{S} = \mathcal{L}$, the case $\mathcal{S} = \mathcal{R}$ is similar. In the base case (rule outIn), we have $P = \overline{a}\langle P''\rangle$ and $R \xrightarrow{\bullet,a,\mathcal{S},P'',\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}} P'$, which implies $P' = \mathbb{E}[\mathbb{F}_2[\mathbb{F}_1[\mathbf{0}] \parallel \mathbb{G}[R''\{P''/X\}]]]$ for some R'' by the first item.

Suppose we are in the case of rule outNu, and let $\mathbb{F}$ such that $\mathsf{extr}(\mathbb{F}) = (\mathbb{F}_1, \mathbb{F}_2)$. Then $P = va.P''$ and $P'' \xrightarrow{\mathbb{F}_1,va.\mathbb{F}_2,\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$. By induction, for all $\mathbb{F}'$ such that $\mathsf{extr}(\mathbb{F}') = (\mathbb{F}_1, va :: \mathbb{F}_2)$, we have $\mathbb{E}[\mathbb{F}'[P''] \parallel R] \to_{\mathsf{rs}} P'$. But since $\mathsf{extr}(\mathbb{F}) = (\mathbb{F}_1, \mathbb{F}_2)$, we have $\mathsf{extr}(va :: \mathbb{F}) = (\mathbb{F}_1, va :: \mathbb{F}_2)$ by definition, so we can apply the induction hypothesis to obtain $\mathbb{E}[(va.\mathbb{F})[P''] \parallel R] \to_{\mathsf{rs}} P'$. This is the same as $\mathbb{E}[\mathbb{F}[va.P''] \parallel R] \to_{\mathsf{rs}} P'$, but $va.P'' = P$, so we get the expected result. The cases of rules outParL and outParR are similar. □

Theorem C.2. *For all $P \to_{\mathsf{zs}} P'$, we have $P \to_{\mathsf{rs}} P'$.*

The proof of the reverse implication follows the same strategy as in HOcore, using the following result.

Lemma C.3. *For all $R \xrightarrow{\mathbb{G},\mathcal{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}} R'$, we have $\mathbb{G}[R] \xrightarrow{\bullet,\mathcal{S},a,P,\mathbb{E},\mathbb{F}_1,\mathbb{F}_2}_{\mathsf{in}} R'$.*
For all $P \xrightarrow{\mathbb{F}_1,\mathbb{F}_2,\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$ and $\mathbb{F}$ such that $\mathsf{extr}(\mathbb{F}) = (\mathbb{F}_1, \mathbb{F}_2)$, we have $\mathbb{F}[P] \xrightarrow{\bullet,\bullet,\mathcal{S},\mathbb{E},R}_{\mathsf{out}} P'$.
For all $P \xrightarrow{\mathbb{E}}_{\mathsf{par}} P'$, we have $\mathbb{E}[P] \xrightarrow{\bullet}_{\mathsf{par}} P'$.

Theorem C.4. *For all $P \to_{\mathsf{rs}} P'$, we have $P \to_{\mathsf{zs}} P'$.*

The NDAM can be found in Figure 9.

$$\langle P \rangle_{\mathsf{zs}} \mapsto \langle P \,;\, \mathsf{init} \mid \bullet \rangle_{\mathsf{par}}$$

$$\begin{aligned}
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} &\mapsto \langle P \,;\, (\mathsf{parL}, \Sigma) :: \pi \mid \parallel Q :: \mathbb{E} \rangle_{\mathsf{par}} && \text{if } \mathsf{par} \notin \mathsf{an}(P) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} &\mapsto \langle Q \,;\, (\mathsf{parR}, \Sigma) :: \pi \mid P \parallel :: \mathbb{E} \rangle_{\mathsf{par}} && \text{if } \mathsf{par} \notin \mathsf{an}(Q) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} &\mapsto \langle P \,;\, (\mathsf{parCom}, \Sigma) :: \pi \mid \bullet, \mathbb{E}, Q \rangle_{\mathsf{left}} && \text{if } (\mathsf{left}, |Q|) \notin \mathsf{an}(P) \\
\langle P \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} &\mapsto \langle \pi \,;\, P^{\cup \mathsf{par}} \mid \mathbb{E} \rangle_{\mathsf{bpar}} && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\langle \mathsf{init} \,;\, P \mid \bullet \rangle_{\mathsf{bpar}} &\mapsto \langle P \rangle_{\mathsf{nf}} \\
\langle (\mathsf{parL}, \Sigma) :: \pi \,;\, P \mid \parallel Q :: \mathbb{E} \rangle_{\mathsf{bpar}} &\mapsto \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} \\
\langle (\mathsf{parR}, \Sigma) :: \pi \,;\, Q \mid P \parallel :: \mathbb{E} \rangle_{\mathsf{bpar}} &\mapsto \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}}
\end{aligned}$$

$$\begin{aligned}
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} &\mapsto \langle P \,;\, (\mathsf{leftParL}, \Sigma) :: \pi \mid \parallel Q :: \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} && \text{if } (\mathsf{left}, |R|) \notin \mathsf{an}(P) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} &\mapsto \langle Q \,;\, (\mathsf{leftParR}, \Sigma) :: \pi \mid P \parallel :: \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} && \text{if } (\mathsf{left}, |R|) \notin \mathsf{an}(Q) \\
\langle \overline{a}^{\Sigma} \langle P \rangle \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} &\mapsto \langle R \,;\, (\mathsf{leftOut}, \Sigma) :: \pi \mid \bullet, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} && \text{if } (\mathsf{in}, a) \notin \mathsf{an}(R) \\
\langle a^{\Sigma}(X).P \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} &\mapsto \langle R \,;\, (\mathsf{leftIn}, \Sigma) :: \pi \mid \bullet, a, \mathbb{F}, \mathbb{E}, X, P \rangle_{\mathsf{out}} && \text{if } (\mathsf{out}, a) \notin \mathsf{an}(R) \\
\langle P \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} &\mapsto \langle \pi \,;\, P^{\cup (\mathsf{left}, R)} \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{bleft}} && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\langle (\mathsf{parCom}, \Sigma) :: \pi \,;\, P \mid \bullet, \mathbb{E}, Q \rangle_{\mathsf{bleft}} &\mapsto \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} \\
\langle (\mathsf{leftParL}, \Sigma) :: \pi \,;\, P \mid \parallel Q :: \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{bleft}} &\mapsto \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} \\
\langle (\mathsf{leftParR}, \Sigma) :: \pi \,;\, Q \mid P \parallel :: \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{bleft}} &\mapsto \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}}
\end{aligned}$$

$$\begin{aligned}
\langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} &\mapsto \langle R \,;\, (\mathsf{inParL}, \Sigma) :: \pi \mid \parallel Q :: \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} && \text{if } (\mathsf{in}, a) \notin \mathsf{an}(R) \\
\langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} &\mapsto \langle Q \,;\, (\mathsf{inParR}, \Sigma) :: \pi \mid R \parallel :: \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} && \text{if } (\mathsf{in}, a) \notin \mathsf{an}(Q) \\
\langle a^{\Sigma}(X).R \,;\, \pi \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} &\mapsto \langle |\mathbb{E}[\mathbb{F}[\mathbf{0}] \parallel \mathbb{G}[R\{P/X\}]]| \rangle_{\mathsf{zs}} \\
\langle R \,;\, \pi \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} &\mapsto \langle \pi \,;\, R^{\cup (\mathsf{in}, a)} \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{bin}} && \text{otherwise}
\end{aligned}$$

$$\begin{aligned}
\langle (\mathsf{leftOut}, \Sigma) :: \pi \,;\, R \mid \bullet, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{bin}} &\mapsto \langle \overline{a}^{\Sigma} \langle P \rangle \,;\, \pi \mid \mathbb{F}, \mathbb{E}, R \rangle_{\mathsf{left}} \\
\langle (\mathsf{inParL}, \Sigma) :: \pi \,;\, R \mid \parallel Q :: \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{bin}} &\mapsto \langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}} \\
\langle (\mathsf{inParR}, \Sigma) :: \pi \,;\, Q \mid R \parallel :: \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{bin}} &\mapsto \langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, a, P, \mathbb{E}, \mathbb{F} \rangle_{\mathsf{in}}
\end{aligned}$$

Figure 8. Left-first NDAM for HOcore

D Correspondence Proofs

Derivations. Let $(\mathcal{S}, \mathcal{O}, \mathcal{R})$ be a zipper semantics with annotation function ϕ. We inductively define that a sequence of rules $\rho_1, \ldots, \rho_n$ is a derivation of a statement as follows. Given an axiom $\rho = \dfrac{\mathcal{P}(\widetilde{w})}{e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} e_t'}$ and a grounding substitution σ which satisfies $\mathcal{P}$, we write $\rho \vdash (e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} e_t')\sigma$. Given a (potentially initial) rule $\rho_1 = \dfrac{e_t' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$ and a grounding substitution σ which satisfies $\mathcal{P}$, we write $\rho_1, \rho_2, \ldots \rho_n \vdash (e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t)\sigma$ if $\rho_2, \ldots \rho_n \vdash (e_t' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t)\sigma$.

$$\langle P \rangle_{\mathsf{zs}} \mapsto \langle P \,;\, \mathsf{init} \mid \bullet \rangle_{\mathsf{par}}$$

$$\begin{array}{rcll}
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} & \mapsto & \langle P \,;\, (\mathsf{parL}, \Sigma) :: \pi \mid \parallel Q :: \mathbb{E} \rangle_{\mathsf{par}} & \text{if } \mathsf{par} \notin \mathsf{an}(P) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} & \mapsto & \langle Q \,;\, (\mathsf{parR}, \Sigma) :: \pi \mid P \parallel :: \mathbb{E} \rangle_{\mathsf{par}} & \text{if } \mathsf{par} \notin \mathsf{an}(Q) \\
\langle \nu^{\Sigma} a.P \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} & \mapsto & \langle P \,;\, (\mathsf{parNu}, \Sigma) :: \pi \mid \nu a :: \mathbb{E} \rangle_{\mathsf{par}} & \text{if } \mathsf{par} \notin \mathsf{an}(P) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} & \mapsto & \langle P \,;\, (\mathsf{parOutL}, \Sigma) :: \pi \mid \bullet, \bullet, \mathcal{L}, \mathbb{E}, Q \rangle_{\mathsf{out}} & \text{if } (\mathsf{out}, |Q|, \bullet) \notin \mathsf{an}(P) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} & \mapsto & \langle Q \,;\, (\mathsf{parOutR}, \Sigma) :: \pi \mid \bullet, \bullet, \mathcal{R}, \mathbb{E}, P \rangle_{\mathsf{out}} & \text{if } (\mathsf{out}, |P|, \bullet) \notin \mathsf{an}(Q) \\
\langle P \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} & \mapsto & \langle \pi \,;\, P^{\cup \mathsf{par}} \mid \mathbb{E} \rangle_{\mathsf{bpar}} & \text{otherwise}
\end{array}$$

$$\begin{array}{rcl}
\langle \mathsf{init} \,;\, P \mid \bullet \rangle_{\mathsf{bpar}} & \mapsto & \langle P \rangle_{\mathsf{nf}} \\
\langle (\mathsf{parL}, \Sigma) :: \pi \,;\, P \mid \parallel Q :: \mathbb{E} \rangle_{\mathsf{bpar}} & \mapsto & \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} \\
\langle (\mathsf{parR}, \Sigma) :: \pi \,;\, Q \mid P \parallel :: \mathbb{E} \rangle_{\mathsf{bpar}} & \mapsto & \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} \\
\langle (\mathsf{parNu}, \Sigma) :: \pi \,;\, P \mid \nu a :: \mathbb{E} \rangle_{\mathsf{bpar}} & \mapsto & \langle \nu^{\Sigma} a.P \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}}
\end{array}$$

$$\begin{array}{rcll}
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \mapsto & \langle P \,;\, (\mathsf{outParL}, \Sigma) :: \pi \mid \parallel Q :: \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \text{if } (\mathsf{out}, |R|, \mathbb{F}_2) \notin \mathsf{an}(P) \\
\langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \mapsto & \langle Q \,;\, (\mathsf{outParR}, \Sigma) :: \pi \mid P \parallel :: \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \text{if } (\mathsf{out}, |R|, \mathbb{F}_2) \notin \mathsf{an}(Q) \\
\langle \nu^{\Sigma} a.P \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \mapsto & \langle P \,;\, (\mathsf{outNu}, \Sigma) :: \pi \mid \mathbb{F}_1, \nu a :: \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \text{if } (\mathsf{out}, |R|, \nu a.\mathbb{F}_2) \notin \mathsf{an}(P) \\
\langle \overline{a}^{\Sigma} \langle P \rangle \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \mapsto & \langle R \,;\, (\mathsf{outIn}, \Sigma) :: \pi \mid \bullet, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \text{if } (\mathsf{in}, a) \notin \mathsf{an}(R), a \notin \mathsf{bn}(\mathbb{F}_2) \\
\langle P \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} & \mapsto & \langle \pi \,;\, P^{\cup (\mathsf{out}, |R|)} \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{bout}} & \text{otherwise}
\end{array}$$

$$\begin{array}{rcl}
\langle (\mathsf{parOutL}, \Sigma) :: \pi \,;\, P \mid \bullet, \bullet, \mathcal{L}, \mathbb{E}, Q \rangle_{\mathsf{bout}} & \mapsto & \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} \\
\langle (\mathsf{parOutR}, \Sigma) :: \pi \,;\, Q \mid \bullet, \bullet, \mathcal{R}, \mathbb{E}, P \rangle_{\mathsf{bout}} & \mapsto & \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{E} \rangle_{\mathsf{par}} \\
\langle (\mathsf{outParL}, \Sigma) :: \pi \,;\, P \mid \parallel Q :: \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{bout}} & \mapsto & \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} \\
\langle (\mathsf{outParR}, \Sigma) :: \pi \,;\, Q \mid P \parallel :: \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{bout}} & \mapsto & \langle P \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} \\
\langle (\mathsf{outNu}, \Sigma) :: \pi \,;\, P \mid \mathbb{F}_1, \nu a :: \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{bout}} & \mapsto & \langle \nu^{\Sigma} a.P \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}}
\end{array}$$

$$\begin{array}{rcll}
\langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \mapsto & \langle R \,;\, (\mathsf{inParL}, \Sigma) :: \pi \mid \parallel Q :: \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \text{if } (\mathsf{in}, a) \notin \mathsf{an}(R) \\
\langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \mapsto & \langle Q \,;\, (\mathsf{inParR}, \Sigma) :: \pi \mid R \parallel :: \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \text{if } (\mathsf{in}, a) \notin \mathsf{an}(Q) \\
\langle \nu^{\Sigma} b.R \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \mapsto & \langle R \,;\, (\mathsf{inNu}, \Sigma) :: \pi \mid \nu b :: \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \text{if } (\mathsf{in}, a) \notin \mathsf{an}(R), b \neq a \\
\langle b^{\Sigma}(X).R \,;\, \pi \mid \mathbb{G}, \mathcal{L}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \mapsto & \langle |\mathbb{E}[\mathbb{F}_2[\mathbb{F}_1[\mathbf{0}] \parallel \mathbb{G}[R\{P/X\}]]]| \rangle_{\mathsf{zs}} & \text{if } a = b \\
\langle b^{\Sigma}(X).R \,;\, \pi \mid \mathbb{G}, \mathcal{R}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \mapsto & \langle |\mathbb{E}[\mathbb{F}_2[\mathbb{G}[R\{P/X\}] \parallel \mathbb{F}_1[\mathbf{0}]]]| \rangle_{\mathsf{zs}} & \text{if } a = b \\
\langle R \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} & \mapsto & \langle \pi \,;\, R^{\cup (\mathsf{in}, a)} \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{bin}} & \text{otherwise}
\end{array}$$

$$\begin{array}{rcl}
\langle (\mathsf{outIn}, \Sigma) :: \pi \,;\, R \mid \bullet, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{bin}} & \mapsto & \langle \overline{a}^{\Sigma} \langle P \rangle \,;\, \pi \mid \mathbb{F}_1, \mathbb{F}_2, \mathcal{S}, \mathbb{E}, R \rangle_{\mathsf{out}} \\
\langle (\mathsf{inParL}, \Sigma) :: \pi \,;\, R \mid \parallel Q :: \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{bin}} & \mapsto & \langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} \\
\langle (\mathsf{inParR}, \Sigma) :: \pi \,;\, Q \mid R \parallel :: \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{bin}} & \mapsto & \langle R \parallel^{\Sigma} Q \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}} \\
\langle (\mathsf{inNu}, \Sigma) :: \pi \,;\, R \mid \nu b :: \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{bin}} & \mapsto & \langle \nu^{\Sigma} b.R \,;\, \pi \mid \mathbb{G}, \mathcal{S}, a, P, \mathbb{E}, \mathbb{F}_1, \mathbb{F}_2 \rangle_{\mathsf{in}}
\end{array}$$

Figure 9. Non-Deterministic Abstract Machine for HOπ

Given a ground statement $e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} e'_t$, we write $\vdash e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} e'_t$ if there exist $\rho_1, \dots \rho_n$ such that $\rho_1, \dots, \rho_n \vdash e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} e'_t$.

Annotations. Consider the HOcore process $R \parallel \overline{a}\langle P\rangle \parallel \overline{a}\langle Q\rangle$ where R cannot do an input on a. A machine run may try first R with an input transition on a with message P, and when it fails to find the input, it annotates R with (in, a). The annotation prevents from testing R for an input on a with message Q, because we know that the success of the input transition does not depend on the message. The annotation contains enough arguments (here a) to know that a term is a normal form w.r.t. any transition with these specific arguments, and independently from the other arguments (like the message).

The next result formalizes the idea that the arguments which are not in the annotation do not matter. It says that if two instances σ and σ' of a transition $\xrightarrow{\widetilde{e}}_{\mathsf{m}}$ agrees on the annotation $\widetilde{f} \subseteq \widetilde{e}$, then a term is a normal form w.r.t. $\xrightarrow{\widetilde{e}\sigma}_{\mathsf{m}}$ iff it is a normal form w.r.t. $\xrightarrow{\widetilde{e}\sigma'}_{\mathsf{m}}$, even if σ and σ' differ on $\widetilde{e} \setminus \widetilde{f}$.

Lemma D.1. *Let* m *be a mode with arguments* $\widetilde{e}$*, and suppose* $\phi(\mathsf{m}, \widetilde{e}) = (\mathsf{m}, \widetilde{f})$*. For all ground term* T *and grounding substitutions* σ, σ' *such that* $\widetilde{f}\sigma = \widetilde{f}\sigma'$ *, we have* $\neg(T \xrightarrow{\widetilde{e}\sigma}_{\mathsf{m}})$ *iff* $\neg(T \xrightarrow{\widetilde{e}\sigma'}_{\mathsf{m}})$*.*

Proof of Lemma D.1. We proceed by induction on the metric of the zipper semantics. We remind that the annotation contains the arguments which appear either in a side-condition or a premise of a rule (cf. Section 5.3). Suppose $\neg(T \xrightarrow{\widetilde{e}\sigma}_{\mathsf{m}})$, the proof is the same in the other direction.

If $\neg(T \xrightarrow{\widetilde{e}\sigma}_{\mathsf{m}})$ because the root operator of T is not parsed in m, then we also have $\neg(T \xrightarrow{\widetilde{e}\sigma'}_{\mathsf{m}})$ for the same reason. Otherwise, we have rules ρ_i parsing the operator which do not apply, either because the side condition or the premise is not satisfied for each of them. The side conditions are not satisfied also with σ', because σ and σ' agree on the annotation, which contains the variables of the side conditions.

Let e^i be the term at the source of the premise for each ρ_i. We necessarily have $e^i\sigma = e^i\sigma'$, because the variables of the arguments which occurs in e^i are in the annotation, and σ and σ' agree on the annotation. The other variables are from the term T itself, because the rules are constructive and reversible. Because $e^i\sigma = e^i\sigma'$, we can apply the induction hypothesis on each of the premises, which therefore do not hold with σ'. In the end, the rule ρ_i also fails with σ' and we have $\neg(T \xrightarrow{\widetilde{e}\sigma'}_{\mathsf{m}})$. □

Lemma 6.2. *Let* $\mathcal{B} = \langle \pi \,;\, A \mid \widetilde{E}\rangle_{\mathsf{bm}}$ *be a valid configuration. We have* $\neg(|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}})$*.*

Proof of Lemma 6.2. Because $\mathcal{B}$ is valid, there exists $\mathcal{I}$ with empty annotation sets such that $\mathcal{I} \mapsto^* \mathcal{B}$. The proof is by induction on the number of machine steps. There are two kinds of transition leading to a backward configuration. The first possibility is that the root operator of A is not pattern-matched in the mode m. In that case, we have directly $\neg(|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}})$.

In the second case, we have rules $\rho_i = \dfrac{e_t^i \xrightarrow{\widetilde{f_i}}_{\mathsf{m}_i} v_t \qquad \mathcal{P}_i(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$ parsing the root operator op of A, and the machine has taken the default step where none of the rules apply. Let σ be a grounding substitution such that $op(v_\Sigma, \widetilde{v})\sigma = A$ and $\widetilde{e}\sigma = \widetilde{E}$. If none of the rules applies because the predicates are not satisfied, then we have directly $\neg(|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}})$.

Otherwise, we have $\phi(\mathsf{m}_i, \widetilde{f_i}\sigma) \in \mathsf{an}(e_t^i\sigma)$ for some rules. Because we start from $\mathcal{I}$ with empty annotation sets, for each of such rules, the annotation has been added in the machine run before getting to $\mathcal{B}$: there exist $\mathcal{B}_i$ such that $\mathcal{I} \mapsto^* \mathcal{B}_i \mapsto^* \mathcal{B}$, $\mathcal{B}_i = \langle \pi_i \,;\, e_t^i\sigma \mid f_i\sigma_i\rangle_{\mathsf{bm}_i}$, and $\phi(\mathsf{m}_i, \widetilde{f_i}\sigma_i) = \phi(\mathsf{m}_i, \widetilde{f_i}\sigma)$. By induction, we have $\neg(|e_t^i\sigma| \xrightarrow{\widetilde{f_i\sigma_i}}_{\mathsf{m}_i})$, which implies $\neg(|e_t^i\sigma| \xrightarrow{\widetilde{f_i\sigma}}_{\mathsf{m}_i})$ by Lemma D.1. As a result, the premises of the rules do not hold, so none of the rules themselves applies to A. We have $\neg(|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}})$ as required. □

Lemma D.2. *For all* $\mathcal{I}$*, there exists a finite set of annotations such that for all search path* $\mathcal{I} \mapsto^* \mathcal{C}$*, the annotations occurring in* $\mathcal{C}$ *are in this set.*

Proof of Lemma D.2. A *partial transition*, denoted by p, is a transition without its result of the form $T \xrightarrow{\widetilde{E}}_{\mathsf{m}}$. We define the set S_T of partial transitions generated from T as follows:

- $T \rightarrow_{\mathsf{zs}} \in S_T$;

- for any $p \in S$, for any rule $\rho = \dfrac{e'_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$, for any grounding σ satisfying $\mathcal{P}$ such that $(e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}})\sigma = p$, we have $(e'_t \xrightarrow{\widetilde{f}}_{\mathsf{m}'})\sigma \in S_T$.

The machine explores S_T until it reaches an axiom finishing the transition, and the annotations generated during the search are built from the the partial transitions in S_T (cf. Section 5.3).

Let $\mathcal{I} = \langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}}$. We prove that S_T is finite. At each step of the process, we add a finite number of partial transitions, because the number of rules is finite and the number of suitable σ is finite when we restrict them to $\mathsf{rv}(e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}})$. For each of such σ corresponds exactly one added premise, because the semantics is machine constructive (cf. Definition 5.1). The process itself cannot go on indefinitely, because the premises are strictly smaller than the conclusion according to the well-foundedness hypothesis (Definition 5.3). □

Theorem 6.3. *For all $\mathcal{I}$, there exists a number n such that any search path starting from $\mathcal{I}$ has length at most n.*

Proof of Theorem 6.3. Assume there exists an initial configuration $\mathcal{I}$ with an infinite search path $\mathcal{I} \mapsto^* \ldots$. Then, by Lemma D.8 and by machine construction there exists an infinite sequence of backtrack-free steps: $\mathcal{I} \overset{\rho}{\longmapsto}_{\mathsf{bf}} \mathcal{F}_1 \overset{\lambda_1}{\longmapsto}_{\mathsf{bf}} \mathcal{F}_2 \overset{\lambda_2}{\longmapsto}_{\mathsf{bf}} \ldots$, where each λ_i is either a τ- or a ρ-label. The number of τ-steps in the sequence is bounded, as each τ-step increases the number of annotations by at least 1, and the number of annotations in total is bounded by the size of the set of Lemma D.2. Ignoring the τ-steps, we are left with an infinite sequence of ρ-steps, which is impossible by the well-foundedness assumption of the zipper semantics underlying the machine construction. □

We show that any machine run eventually reaches an initial or normal-form configuration. If it is possible to get to any of the reducts of a term T when we start from $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}}$ (Theorem 6.5), it is no longer the case once the machine has done some non-deterministic choice. For instance in HOcore, if the machine starts exploring $P \parallel Q$ with P which contains a redex, then Q will be ignored. The machine backtracks only to undo choices that lead to a dead-end, not the ones that eventually lead to a redex.

Given a partial run $\mathcal{F} \mapsto^* \mathcal{F}'$, either we can reach a redex from $\mathcal{F}'$, or we need to backtrack. In the latter case, the machine backtracks as little as possible, i.e., to the first configuration from which we can find a redex. We formalize this idea in the next lemma, using backtrack-free closure.

Lemma D.3. *Let $\mathcal{F} = \langle A \,;\, \pi \mid \widetilde{E} \rangle_{\mathsf{m}}$ be a valid configuration such that $\mathcal{F} \overset{\rho_1 \ldots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{F}'$ and $|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}} T'$ for some T'.*

There exists $0 \le k \le n$ such that for all T' and rules $\rho'_{k+1} \ldots \rho'_m$ verifying $\rho_1 \ldots \rho_k, \rho'_{k+1} \ldots \rho_m \vdash |A| \xrightarrow{\widetilde{E}}_m T'$, there exists A' such that $\mathcal{F} \overset{\rho_1 \ldots \rho_k, \rho'_{k+1} \cdots \rho_m}{\longmapsto}_{\mathsf{bf}} \langle A' \rangle_{\mathsf{zs}}$ and $|A'| = T'$. Besides, for all $k < k' \le n$, for all T', $\rho'_{k'+1} \ldots \rho'_m$, we have $\neg(\rho_1 \ldots \rho_{k'} \rho'_{k'+1} \ldots \rho'_m \vdash |A| \xrightarrow{\widetilde{E}}_{\mathsf{m}} T')$.

Proof of Lemma D.3. Consider all the derivations proving a transition $\vdash |A| \xrightarrow{\widetilde{E}}_{\mathsf{m}} T'$, and take k as the length of the longest common prefix of $\rho_1 \ldots \rho_n$ with these derivations—we may have $k = 0$ if a bad choice is made from the start. $\mathcal{F}_k = \langle A_k \,;\, \pi_k \mid \widetilde{f_k} \rangle_{\mathsf{m}_k}$ be the configuration such that $\mathcal{F} \overset{\rho_1 \ldots \rho_k}{\longmapsto}_{\mathsf{bf}} \mathcal{F}_k$. We show that the machine can backtrack from $\mathcal{F}'$ to a configuration equal to $\mathcal{F}_k$ up to annotations.

If $k = n$, then $\mathcal{F}_k = \mathcal{F}'$ and we do not need to backtrack. Suppose $k < n$, and let $\mathcal{F}_n = \langle A_n \,;\, \pi_n \mid \widetilde{f_n} \rangle_{\mathsf{m}_n}$. We necessarily have $\neg(|A_n| \xrightarrow{\widetilde{f_n}}_{\mathsf{m}_n})$, otherwise we would have $k = n$. Therefore, by Lemma D.5, there exists A'_n such that $\mathcal{F}_n \mapsto^* \langle \pi_n \,;\, A'_n \mid \widetilde{f_n} \rangle_{\mathsf{bm}_n}$ and $|A'_n| = |A_n|$. This configuration then backtracks to a configuration equal to $\mathcal{F}_{n-1}$ up to annotations. By induction, we show that we can reach $\mathcal{F}_k$. Then from this configuration, we can reach any T' verifying $\vdash |A_k| \xrightarrow{\widetilde{f_k}}_{\mathsf{m}_k} T'$ by Lemma D.4, hence the result holds. □

Theorem 6.4. *Let $\mathcal{C}$ be a valid configuration. Either $\mathcal{C} \mapsto^* \mathcal{I}$ for some $\mathcal{I}$, or $\mathcal{C} \mapsto^* \langle A \rangle_{\mathsf{nf}}$ for some A.*

Proof of Theorem 6.4. Suppose $\mathcal{C}$ is a forward configuration $\mathcal{C} = \langle A \,;\, \pi \mid \widetilde{E} \rangle_{\mathsf{m}}$. If $\neg(|A| \xrightarrow{\widetilde{E}}_{\mathsf{m}})$, then Lemma D.5 applies, otherwise we conclude with Lemma D.3.

If $\mathcal{C}$ is a backward configuration, it goes in one step to a forward configuration, and we conclude as in the previous case. □

Lemma D.4. *If $\vdash T \xrightarrow{\widetilde{E}}_{\mathsf{m}} T'$, then for all valid configuration $\mathcal{C} = \langle A ; \pi \mid \widetilde{E} \rangle_{\mathsf{m}}$ such that $|A| = T$, there exists A' such that $\mathcal{C} \mapsto^* \langle A' \rangle_{\mathsf{zs}}$ and $|A'| = T'$.*

Proof of Lemma D.4. By induction on the size of the derivation $\rho_1, \ldots, \rho_n \vdash T \xrightarrow{\widetilde{E}}_{\mathsf{m}} T'$. If $n = 1$, then we apply an axiom, and the corresponding machine step applies directly.

Otherwise, we apply the rule $\rho_1 = \dfrac{e_t'' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t \qquad \mathcal{P}(\widetilde{w})}{e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}$ for some σ such that $(e_t \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t)\sigma = T \xrightarrow{\widetilde{E}}_{\mathsf{m}} T'$. Let $\mathcal{C}' = \langle e_t''\sigma ; (\rho_1, \mathsf{an}(A)) :: \pi \mid \widetilde{f}\sigma \rangle_{\mathsf{m}'}$. We show that we can apply the forward step corresponding to ρ_1, i.e., $\mathcal{C} \mapsto \mathcal{C}'$. The step cannot apply only if the annotation prevents it: suppose we have $\phi(\mathsf{m}', \widetilde{f}\sigma) \in \mathsf{an}(e_t''\sigma)$. Because $\mathcal{C}$ is valid, it is derived from an initial configuration $\mathcal{I}$ without annotations, so the annotation has been added in an intermediary backward configuration: there exist A'', σ', and π' such that $\mathcal{I} \mapsto^* \langle \pi' ; A'' \mid \widetilde{f}\sigma' \rangle_{\mathsf{bm}'} \mapsto^* \mathcal{C}$, $|A''| = e_t''\sigma$, and $\phi(\mathsf{m}', \widetilde{f}\sigma') = \phi(\mathsf{m}', \widetilde{f}\sigma)$. By Lemma 6.2, we have $\neg(|A''| \xrightarrow{\widetilde{f}\sigma'}_{\mathsf{m}'})$, which implies $\neg(|A''| \xrightarrow{\widetilde{f}\sigma}_{\mathsf{m}'})$ by Lemma D.1. We have a contradiction, since $\vdash (e_t'' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t)\sigma$ holds. Therefore we have $\phi(\mathsf{m}', \widetilde{f}\sigma) \notin \mathsf{an}(e_t''\sigma)$ and $\mathcal{C} \mapsto \mathcal{C}'$.

We also have $\rho_2, \ldots, \rho_n \vdash (e_t'' \xrightarrow{\widetilde{f}}_{\mathsf{m}'} v_t)\sigma$, so by induction there exists A' such that $\mathcal{C}' \mapsto^* \langle A' \rangle_{\mathsf{zs}}$ and $|A'| = T'$, from which we can conclude the proof. □

Theorem 6.5. *For all $\vdash T \rightarrow_{\mathsf{zs}} T'$, we have $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}} \mapsto^* \langle \|T'\|^{\emptyset} \rangle_{\mathsf{zs}}$.*

The following lemma expresses the idea that if a term cannot reduce in a given mode, any path from a forward configuration corresponding to these term and mode goes through a backward configuration of this term and mode.

Lemma D.5. *If $\neg(T \xrightarrow{\widetilde{E}}_{\mathsf{m}})$, then for all valid configuration $\mathcal{C} = \langle A ; \pi \mid \widetilde{E} \rangle_{\mathsf{m}}$ such that $|A| = T$, for all search path $\mathcal{C} \mapsto^* \mathcal{C}'$, there exists A' such that either $\mathcal{C}' \mapsto^* \langle \pi ; A' \mid \widetilde{E} \rangle_{\mathsf{bm}}$ or $\langle \pi ; A' \mid \widetilde{E} \rangle_{\mathsf{bm}} \mapsto^* \mathcal{C}'$ with $|A'| = T$.*

Proof of Lemma D.5. We proceed by induction on the metric of the zipper semantics, and distinguish two cases. If the root operator of the zipper semantics is not parsed in m, then the only machine step possible from $\mathcal{C}$ is the step to $\langle \pi ; A' \mid \widetilde{E} \rangle_{\mathsf{bm}}$, where A' is A where the annotation set of its root operator is extended with $\phi(\mathsf{m}, \widetilde{E})$, so the result holds.

Otherwise, there exist rules $\dfrac{e_t^i \xrightarrow{\widetilde{f}_i}_{\mathsf{m}_i} v_t \qquad \mathcal{P}_i(\widetilde{w})}{op(\widetilde{v}) \xrightarrow{\widetilde{e}}_{\mathsf{m}} v_t}\ \rho_i$ parsing the root operator op of T. Let n be the number of rules for which $\mathcal{P}_i$ is satisfied and $\phi(\mathsf{m}_i, \widetilde{f}_i\sigma) \notin \mathsf{an}(e_t^i\sigma)$ for σ a grounding substitution such that $op(v_\Sigma, \widetilde{v}) = A$ and $\widetilde{e}\sigma = \widetilde{E}$; this is the number of forward configurations $\mathcal{C}'$ than can be reached in one step from $\mathcal{C}$. We prove the result by an inner induction on n. If $n = 0$, then no forward step is possible, and only the step to $\langle \pi ; A' \mid \widetilde{E} \rangle_{\mathsf{bm}}$ can be done, where A' is as above.

Otherwise, we can make a step $\langle A ; \pi \mid \widetilde{E} \rangle_{\mathsf{m}} \mapsto \langle e_t^j\sigma ; (\rho_j, \mathsf{an}(A)) :: \pi \mid \widetilde{f}_j\sigma \rangle_{\mathsf{m}_j'}$ for some j. However, we have $\neg(|e_t^j\sigma| \xrightarrow{\widetilde{f}_j\sigma}_{\mathsf{m}'})$, otherwise T could do a m-transition. Because $\xrightarrow{f_j}_{\mathsf{m}'}$ is smaller than $\xrightarrow{\widetilde{e}}_{\mathsf{m}}$ according to the zipper semantics metric, we can apply the outermost induction hypothesis. There exists A' such that $\langle e_t^j\sigma ; (\rho_j, \mathsf{an}(A)) :: \pi \mid \widetilde{f}_j\sigma \rangle_{\mathsf{m}_j'} \mapsto^* \langle (\rho_j, \mathsf{an}(A)) :: \pi ; A' \mid \widetilde{f}_j\sigma \rangle_{\mathsf{bm}_j'}$ and $|A'| = |e_t^j\sigma|$. By construction, the only possible next step is $\langle (\rho_j, \mathsf{an}(A)) :: \pi ; A' \mid \widetilde{f}_j\sigma \rangle_{\mathsf{bm}_j'} \mapsto \langle A'' ; \pi \mid \widetilde{E} \rangle_{\mathsf{m}}$, where A'' differs from A only in their annotation sets. In particular, A'' can no longer do the step corresponding to ρ_j, so $n - 1$ configurations are reachable from A''. Therefore, we can conclude using the induction hypothesis on n. □

As a result, the only possible outcome for a machine run starting from a normal form is a normal-form configuration which cannot step further.

Theorem 6.6. *If $\neg(T \rightarrow_{\mathsf{zs}})$, then any machine run from the configuration $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}}$ ends with $\langle A \rangle_{\mathsf{nf}}$ such that $|A| = T$.*

Lemma D.6. *A step $\overset{\rho}{\mapsto}_{\mathsf{bf}}$ is between two forward configurations, a step $\overset{-\rho}{\longmapsto}_{\mathsf{bf}}$ is between a backward and a forward configuration, and $\overset{\tau}{\mapsto}_{\mathsf{bf}}$ is either between a forward and a backward configuration, or two forward configurations.*

Proof of Lemma D.6. By induction on the derivation of $\overset{\lambda}{\mapsto}_{\mathsf{bf}}$. □

Given two configuration $\mathcal{C}_1$ and $\mathcal{C}_2$, we write $|\mathcal{C}_1| = |\mathcal{C}_2|$ if they are equal up to their terms under focus A_1 and A_2, for which we have $|A_1| = |A_2|$.

Lemma D.7. *For all $\mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{F}'$, we have $|\mathcal{F}| = |\mathcal{F}'|$. For any other transition $\mathcal{C} \overset{\lambda}{\mapsto}_{\mathsf{bf}} \mathcal{C}'$, there exists $\mathcal{C}''$ such that $\mathcal{C} \overset{\lambda}{\mapsto} \mathcal{C}''$ and $|\mathcal{C}''| = |\mathcal{C}'|$.*

Proof of Lemma D.7. We proceed by induction on the derivation of $\mathcal{C} \overset{\lambda}{\mapsto}_{\mathsf{bf}} \mathcal{C}'$. The base case is straightforward. If $\mathcal{F} \overset{\rho}{\mapsto}_{\mathsf{bf}} \overset{\tau}{\mapsto}_{\mathsf{bf}} \overset{-\rho}{\longmapsto}_{\mathsf{bf}} \mathcal{F}'$, then by the induction hypothesis, there exists $\mathcal{F}_1$, $\mathcal{B}_2$, and $\mathcal{C}''$ such that $\mathcal{F} \overset{\rho}{\mapsto} \mathcal{F}_1 \overset{\tau}{\mapsto} \mathcal{B}_2 \overset{-\rho}{\longmapsto} \mathcal{C}''$ and $|\mathcal{C}''| = |\mathcal{F}'|$. The first step applies ρ which is then unapplied, so one can check that $|\mathcal{C}''| = |\mathcal{F}|$, and therefore $|\mathcal{F}| = |\mathcal{F}'|$ as required.

The case $\mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{F}'$ is easy by induction. If $\mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{F}_0 \overset{\rho}{\mapsto}_{\mathsf{bf}} \mathcal{F}'$, then by the induction hypothesis $|\mathcal{F}| = |\mathcal{F}_0|$ and $\mathcal{F}_0 \overset{\rho}{\mapsto} \mathcal{F}''$ for some $\mathcal{F}''$ such that $|\mathcal{F}'| = |\mathcal{F}''|$. If $\mathcal{F}_0$ is able to do a $\overset{\rho}{\mapsto}$ step, then so is $\mathcal{F}$, since $|\mathcal{F}| = |\mathcal{F}_0|$ and $\mathcal{F}_0$ contains more annotations than $\mathcal{F}$. As a result, we have $\mathcal{F} \overset{\rho}{\mapsto} \mathcal{F}''$ with $|\mathcal{F}'| = |\mathcal{F}''|$, as wished. □

For all $\mathcal{C}$, we write $\mathsf{stack}(\mathcal{C})$ for its rules stack.

Lemma D.8. *Let $\mathcal{I}$ be a valid initial configuration. For all $\mathcal{I} \mapsto^* \overset{\rho}{\mapsto} \mathcal{C}$, there exist $\rho_1 \dots \rho_n$ such that $\mathcal{I} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{C}$.*

For all $\mathcal{I} \mapsto^* \overset{\lambda}{\mapsto} \mathcal{C}$ such that $\lambda \neq \rho$, there exist $\mathcal{F}$, $\rho_1 \dots \rho_n$ such that $\mathcal{I} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{C}$ and $\mathsf{stack}(\mathcal{F}) = \mathsf{stack}(\mathcal{C})$.

Proof of Lemma D.8. We proceed by induction on the number of machine steps. The base case is easy. In the inductive case, we distinguish three cases depending on the last step.

Let $\mathcal{I} \mapsto^* \mathcal{F} \overset{\tau}{\mapsto} \mathcal{C}$, and assume the step before $\mathcal{F}$ is $\lambda \neq \rho$. By the induction hypothesis, there exist $\mathcal{F}'$, $\rho_1 \dots \rho_n$ such that $\mathcal{I} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{F}' \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{F}$ and $\mathsf{stack}(\mathcal{F}') = \mathsf{stack}(\mathcal{F})$. By definition of the closure, we have $\mathcal{F}' \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{C}$, and by the shape of switching rules, we also have $\mathsf{stack}(\mathcal{F}) = \mathsf{stack}(\mathcal{C})$, so we can conclude. The case $\lambda = \rho$ is easier.

Let $\mathcal{I} \mapsto^* \mathcal{C}' \overset{\rho}{\mapsto} \mathcal{C}$, and assume the step before $\mathcal{C}'$ is $\lambda \neq \rho'$. By the induction hypothesis, there exist $\mathcal{F}$, $\rho_1 \dots \rho_n$ such that $\mathcal{I} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{C}'$. By definition of the closure, we have $\mathcal{F} \overset{\rho}{\mapsto}_{\mathsf{bf}} \mathcal{C}$, therefore $\mathcal{I} \overset{\rho_1 \dots \rho_n \rho}{\longmapsto}_{\mathsf{bf}} \mathcal{C}$ holds, as wished. The case $\lambda = \rho'$ is easier.

If $\mathcal{I} \mapsto^* \mathcal{B} \overset{-\rho}{\longmapsto} \mathcal{C}$, then the step before $\mathcal{B}$ is necessarily a τ-step. By the induction hypothesis, there exist $\mathcal{F}$, $\rho_1 \dots \rho_n, \rho_{n+1}$ such that $\mathcal{I} \overset{\rho_1 \dots \rho_n \rho_{n+1}}{\longmapsto}_{\mathsf{bf}} \mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{B}$ and $\mathsf{stack}(\mathcal{F}) = \mathsf{stack}(\mathcal{B})$. Let $\mathcal{F}'$ be the configuration such that $\mathcal{I} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{F}' \overset{\rho_{n+1}}{\longmapsto}_{\mathsf{bf}} \mathcal{F}$. We have $(\rho_{n+1}, \Sigma) :: \mathsf{stack}(\mathcal{F}') = \mathsf{stack}(\mathcal{F})$ for some Σ by Lemma D.7 and $\mathsf{stack}(\mathcal{B}) = (\rho, \Sigma') :: \mathsf{stack}(\mathcal{C})$ for some Σ' by construction. Because $\mathsf{stack}(\mathcal{F}) = \mathsf{stack}(\mathcal{B})$, we have $\rho = \rho_{n+1}$ and $\mathsf{stack}(\mathcal{F}') = \mathsf{stack}(\mathcal{C})$. Therefore we have $\mathcal{I} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \mathcal{F}' \overset{\rho}{\mapsto}_{\mathsf{bf}} \mathcal{F} \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{B} \overset{-\rho}{\longmapsto} \mathcal{C}$, i.e., $\mathcal{F}' \overset{\tau}{\mapsto}_{\mathsf{bf}} \mathcal{C}$ by definition of the closure, hence the result holds. □

The condition on stacks in the second case implies that a potential backtracking step after the τ step, it will undo the last rule applied before $\mathcal{F}$, i.e., ρ_n.

Lemma D.9. *Let $\langle A \,;\, \pi \mid \widetilde{E} \rangle_{\mathsf{m}}$ be a valid configuration such that there exist $\rho_1 \dots \rho_n$, A' so that $\langle A \,;\, \pi \mid \widetilde{E} \rangle_{\mathsf{m}} \overset{\rho_1 \dots \rho_n}{\longmapsto}_{\mathsf{bf}} \langle A' \rangle_{\mathsf{zs}}$ is a search path. Then $\rho_1 \dots \rho_n \vdash |A| \overset{\widetilde{E}}{\longrightarrow}_{\mathsf{m}} |A'|$.*

Proof of Lemma D.9. We proceed by induction on n. If $n = 1$, we apply an axiom. By Lemma D.7, we have $\langle A \,;\, \pi \mid \widetilde{E} \rangle_{\mathsf{m}} \overset{\rho_1}{\longmapsto} \langle A'' \rangle_{\mathsf{zs}}$ for some A'' such that $|A''| = |A'|$. An instance of a machine step implies an instance of the axiom, therefore we have $\rho_1 \vdash |A| \overset{\widetilde{E}}{\longrightarrow}_{\mathsf{m}} |A'|$.

If $n > 1$, we apply the induction hypothesis on the sequence $\rho_2 \dots \rho_n$ to get $\rho_2 \dots \rho_n \vdash |A''| \overset{\widetilde{E}}{\longrightarrow}_{\mathsf{m}} |A'|$ for some A''. By Lemma D.7, we have $\langle A \,;\, \pi \mid \widetilde{E} \rangle_{\mathsf{m}} \overset{\rho_1}{\longmapsto} \langle A''' \,;\, (\rho_1, \Sigma) :: \pi \mid \widetilde{F} \rangle_{\mathsf{m}'}$ where A'' is such that $|A'''| = |A''|$. Let σ be the grounding substitution of this machine step, and v_T be the rule variable designating the outcome of the transitions in the source and premise of ρ_1. We have an instance of ρ_1 by considering σ' which maps any $w \neq v_T$ to $|w\sigma|$ and v_T to $|A'|$. This instance completes the derivation and we have $\rho_1 \dots \rho_n \vdash |A| \overset{\widetilde{E}}{\longrightarrow}_{\mathsf{m}} |A'|$, as wished. □

Theorem 6.7. *For all search path $\langle \|T\|^{\emptyset} \rangle_{\mathsf{zs}} \mapsto^+ \langle A \rangle_{\mathsf{zs}}$, we have $\vdash T \rightarrow_{\mathsf{zs}} |A|$.*

Proof of Theorem 6.7. We first apply Lemma D.8; because the machine run ends with an initial configuration, the last step is a forward step (an axiom application), so we are in the first case of the lemma. We can apply Lemma D.9 to the resulting backtrack-free sequence to obtain the zipper semantics derivation. □

Lemma A.1. *For all* $t \xrightarrow{s,\mathbb{E}}_{\mathsf{lam}} t'$, *we have* $\mathbb{E}[t \mathbin{@} s] \to_{\mathsf{rs}} t'$.

For all $t \xrightarrow{\mathbb{E}}_{\mathsf{app}} t'$, *we have* $\mathbb{E}[t] \to_{\mathsf{rs}} t'$.

Theorem A.4. *For all* $t \to_{\mathsf{rs}} t'$, *we have* $t \to_{\mathsf{zs}} t'$.