A Complete Normal-Form Bisimilarity for State

Dariusz Biernacki[1], Sergueï Lenglet[2(⊠)], and Piotr Polesiuk[1]

[1] University of Wrocław, Wrocław, Poland
{dabi,ppolesiuk}@cs.uni.wroc.pl
[2] Université de Lorraine, Nancy, France
serguei.lenglet@univ-lorraine.fr

Abstract. We present a sound and complete bisimilarity for an untyped λ-calculus with higher-order local references. Our relation compares values by applying them to a fresh variable, like normal-form bisimilarity, and it uses environments to account for the evolving store. We achieve completeness by a careful treatment of evaluation contexts comprising open stuck terms. This work improves over Støvring and Lassen's incomplete environment-based normal-form bisimilarity for the $\lambda\rho$-calculus, and confirms, in relatively elementary terms, Jaber and Tabareau's result, that the state construct is discriminative enough to be characterized with a bisimilarity without any quantification over testing arguments.

1 Introduction

Two terms are contextually equivalent if replacing one by the other in a bigger program does not change the behavior of the program. The quantification over program contexts makes contextual equivalence hard to use in practice and it is therefore common to look for more effective characterizations of this relation. In a calculus with local state, such a characterization has been achieved either through *logical relations* [1,5,15], which rely on types, denotational models [6,10,13], or coinductively defined *bisimilarities* [9,12,17–19].

Koutavas et al. [8] argue that to be sound w.r.t. contextual equivalence, a bisimilarity for state should accumulate the tested terms in an environment to be able to try them again as the store evolves. Such *environmental bisimilarities* usually compare terms by applying them to arguments built from the environment [12,17,19], and therefore still rely on some universal quantification over testing arguments. An exception is Støvring and Lassen's bisimilarity [18], which compares terms by applying them to a fresh variable, like one would do with a *normal-form* (or *open*) bisimilarity [11,16]. Their bisimilarity characterizes contextual equivalence in a calculus with control and state, but is not *complete* in a calculus with state only: there exist equivalent terms that are not related by the bisimilarity. Jaber and Tabareau [6] go further and propose a sound and complete *Kripke Open Bisimilarity* for a calculus with local state, which also compares terms by applying them to a fresh variable, but uses notions from Kripke logical relations, namely transition systems of invariants, to reason about heaps.

In this paper, we propose a sound and complete normal-form bisimilarity for a call-by-value λ-calculus with local references which relies on environments to handle heaps. We therefore improve over Støvring and Lassen's work, since our relation is complete, by following a different, potentially simpler, path than Jaber and Tabareau, since we use environments to represent possible worlds and do not rely on any external structures such as transition systems of invariants. Moreover, we do not need types and define our relation in an untyped calculus.

We obtain completeness by treating carefully normal forms that are not values, i.e., open stuck terms of the form $E[x\,v]$. First, we distinguish in the environment the terms which should be tested multiple times from the ones that should be run only once, namely the evaluation contexts like E in the above term. The latter are kept in a separate environment that takes the form of a stack, according to the idea presented by Laird [10] and by Jagadeesan et al. [7]. Second, we relate the so-called *deferred diverging* terms [5,6], i.e., open stuck terms which hide a diverging behavior in the evaluation context E, with the regular diverging terms.

It may be worth stressing that our congruence proof is based on the machinery we have developed before [3] and is simpler than Støvring and Lassen's one, in particular in how it accounts for the extensionality of functions.

We believe that this work makes a contribution to the understanding of how one should adjust the normal-form bisimulation proof principle when the calculus under consideration becomes less discriminative, assuming that one wishes to preserve completeness of the theory. In particular, it is quite straightforward to define a complete normal-form bisimilarity for the λ-calculus with first-class continuations and global store, with no need to refer to other notions than the ones already present in the reduction semantics. Similarly, in the $\lambda\mu\rho$-calculus (continuations and local references), one only needs to introduce environments to ensure soundness of the theory, but essentially nothing more is required to obtain completeness [18]. In this article we show which new ingredients are needed when moving from these two highly expressive calculi to the corresponding, less discriminative ones—with global or local references only—that do not offer access to the current continuation.

The rest of this paper is as follows. In Sect. 2, we study a simple calculus with global store to see how to reach completeness in that case. In particular, we show in Sect. 2.2 how we deal with deferred diverging terms. We remind in Sect. 2.3 the notion of *diacritical progress* [3] and the framework our bisimilarity and its proof of soundness are based upon. We sketch the completeness proof in Sect. 2.4. Section 2 paves the way for the main result of the paper, described in Sect. 3, where we turn to the calculus with local store. We define the bisimilarity in Sect. 3.2, prove its soundness and completeness in Sect. 3.3, and use it in Sect. 3.4 on examples taken from the literature. We conclude in Sect. 4, where we discuss related work and in particular compare our work to Jaber and Tabareau's. A companion report expands on the proofs [4].

2   Global Store

We first consider a calculus where terms share a global store and present how we deal with deferred diverging terms to get a complete bisimilarity.

2.1   Syntax, Semantics, and Contextual Equivalence

We extend the call-by-value λ-calculus with the ability to read and write a global memory. We let x, y, ... range over term variables and l range over references. A *store*, denoted by h, g, is a finite map from references to values; we write $\mathsf{dom}(h)$ for the domain of h, i.e., the set of references on which h is defined. We write $\emptyset$ for the empty store, $h \uplus g$ for the union of two stores, assuming $\mathsf{dom}(h) \cap \mathsf{dom}(g) = \emptyset$. The syntax of terms and contexts is defined as follows.

Terms:	$t, s ::= v \mid t\,t \mid l := t; t \mid\, !l$
Values:	$v, w ::= x \mid \lambda x.t$
Evaluation contexts:	$E, F ::= \Box \mid E\,t \mid v\,E \mid l := E; t$

The term $l := t; s$ evaluates t (if possible) and stores the resulting value in l before continuing as s, while $!l$ reads the value kept in l. When writing examples and in the completeness proofs, we use natural numbers, booleans, the conditional if ... then ... else ..., local definitions let ... in ..., sequence ;, and unit () assuming the usual call-by-value encodings for these constructs.

A λ-abstraction $\lambda x.t$ binds x in t; we write $\mathsf{fv}(t)$ (respectively $\mathsf{fv}(E)$) for the set of free variables of t (respectively E). We identify terms up to α-conversion of their bound variables. A variable or reference is *fresh* if it does not occur in any other entities under consideration, and a store is fresh if it maps references to pairwise distinct fresh variables. A term or context is *closed* if it has no free variables. We write $\mathsf{fr}(t)$ for the set of references that occur in t.

The call-by-value semantics of the calculus is defined on *configurations* $\langle h \mid t \rangle$ such that $\mathsf{fr}(t) \subseteq \mathsf{dom}(h)$ and for all $l \in \mathsf{dom}(h)$, $\mathsf{fr}(h(l)) \subseteq \mathsf{dom}(h)$. We let c and d range over configurations. We write $t\{v/x\}$ for the usual capture-avoiding substitution of x by v in t, and we let $\int$ range over simultaneous substitutions $.\{v_1/x_1\} \ldots \{v_n/x_n\}$. We write $h[l := v]$ for the operation updating the value of l to v. The reduction semantics $\to$ is defined by the following rules.

$$\langle h \mid (\lambda x.t)\,v \rangle \to \langle h \mid t\{v/x\} \rangle \qquad \langle h \mid\, !l \rangle \to \langle h \mid h(l) \rangle$$
$$\langle h \mid l := v; t \rangle \to \langle h[l := v] \mid t \rangle \qquad \langle h \mid E[t] \rangle \to \langle g \mid E[s] \rangle \text{ if } \langle h \mid t \rangle \to \langle g \mid s \rangle$$

The well-formedness condition on configurations ensures that a read operation $!l$ cannot fail. We write $\to^*$ for the reflexive and transitive closure of $\to$.

A term t of a configuration $\langle h \mid t \rangle$ which cannot reduce further is called a *normal form*. Normal forms are either values or *open-stuck terms* of the form $E[x\,v]$; closed normal forms can only be λ-abstractions. A configuration *terminates*, written $c \Downarrow$ if it reduces to a normal-form configuration; otherwise it *diverges*, written $c \Uparrow$, like configurations running $\Omega \overset{\mathsf{def}}{=} (\lambda x.x\,x)\,(\lambda x.x\,x)$.

Contextual equivalence equates terms behaving the same in all contexts. A substitution $\int$ closes a term t if $t\int$ is closed; it closes a configuration $\langle h \mid t \rangle$ if it closes t and the values in h.

Definition 1. *t and s are contextually equivalent, written $t \equiv s$, if for all contexts E, fresh stores h, and closing substitutions $\int$, $\langle h \mid E[t] \rangle \int \Downarrow$ iff $\langle h \mid E[s] \rangle \int \Downarrow$.*

Testing only evaluation contexts is not a restriction, as it implies the equivalence w.r.t. all contexts $\equiv_C$: one can show that $t \equiv_C s$ iff $\lambda x.t \equiv_C \lambda x.s$ iff $\lambda x.t \equiv \lambda x.s$.

2.2   Normal-Form Bisimulation

Informal Presentation. Two open terms are normal-form bisimilar if their normal forms can be decomposed into bisimilar subterms. For example in the plain λ-calculus, a stuck term $E[x\,v]$ is bisimilar to t if t reduces to a stuck term $F[xw]$ so that respectively E, F and v, w are bisimilar when they are respectively plugged with and applied to a fresh variable.

Such a requirement is too discriminating for many languages, as it distinguishes terms that should be equivalent. For instance in plain λ-calculus, given a closed value v, $t \stackrel{\mathsf{def}}{=} x\,v$ is not normal form bisimilar to $s \stackrel{\mathsf{def}}{=} (\lambda y.x\,v)\,(x\,v)$. Indeed, $\square$ is not bisimilar to $(\lambda y.x\,v)\,\square$ when plugged with a fresh z: the former produces a value z while the latter reduces to a stuck term $x\,v$. However, t and s are contextually equivalent, as for all closed value w, $t\{w/x\}$ and $s\{w/x\}$ behave like $w\,v$: if $w\,v$ diverges, then they both diverges, and if $w\,v$ evaluates to some value w', then they also evaluates to w'. Similarly, $x\,v\,\Omega$ and Ω are not normal-form bisimilar (one is a stuck term while the other is diverging), but they are contextually equivalent by the same reasoning.

The terms t and s are no longer contextually equivalent in a λ-calculus with store, since a function can count how many times it is applied and change its behavior accordingly. More precisely, t and s are distinguished by the context $l := 0; (\lambda x.\square)\,\lambda z.l := !l + 1;$ if $!l = 1$ then 0 else Ω. But this counting trick is not enough to discriminate $x\,v\,\Omega$ and Ω, as they are still equivalent in a λ-calculus with store. Although $x\,v\,\Omega$ is a normal form, it is in fact always diverging when we replace x by an arbitrary closed value w, either because $w\,v$ itself diverges, or it evaluates to some w' and then $w'\,\Omega$ diverges. A stuck term which hides a diverging behavior has been called *deferred diverging* in the literature [5,6].

It turns out that being able to relate a diverging term to a deferred diverging term is all we need to change from the plain λ-calculus normal-form bisimilarity to get a complete equivalence when we add global store. We do so by distinguishing two cases in the clause for open-stuck terms: a configuration $\langle h \mid E[x\,v] \rangle$ is related to c either if c can reduce to a stuck configuration with related subterms, or if E is a diverging context, and we do not require anything of c. The resulting simulation is not symmetric as it relates a deferred diverging configuration with any configuration c (even converging one), but the corresponding notion of bisimulation equates such configuration only to either a configuration of the same kind or a diverging configuration such as $\langle h \mid \Omega \rangle$.

Progress. We define simulation using the notion of *diacritical progress* we developed in a previous work [2,3], which distinguishes between *active* and *passive* clauses. Roughly, passive clauses are between simulation states which should be considered equal, while active clauses are between states where actual progress is taking place. This distinction does not change the notions of bisimulation or bisimilarity, but it simplifies the soundness proof of the bisimilarity. It also allows for the definition of powerful *up-to techniques*, relations that are easier to use than bisimulations but still imply bisimilarity. For normal-form bisimilarity, our framework enables up-to techniques which respects η-expansion [3].

Progress is defined between objects called *candidate relations*, denoted by $\mathcal{R}$, $\mathcal{S}$, $\mathcal{T}$. A candidate relation $\mathcal{R}$ contains pairs of configurations, and a set of configurations written $\mathcal{R}{\uparrow}$, which we expect to be composed of diverging or deferred diverging configurations (for such relations we take $\mathcal{R}^{-1}{\uparrow}$ to be $\mathcal{R}{\uparrow}$). We extend $\mathcal{R}$ to stores, terms, values, and contexts with the following definitions.

$$\frac{\mathsf{dom}(h) = \mathsf{dom}(g) \quad \forall l, h(l)\ \mathcal{R}^{\mathsf{v}}\ g(l)}{h\ \mathcal{R}^{\mathsf{h}}\ g} \qquad \frac{\langle h \mid t \rangle\ \mathcal{R}\ \langle h \mid s \rangle \quad h\ \text{fresh}}{t\ \mathcal{R}^{\mathsf{t}}\ s}$$

$$\frac{v\,x\ \mathcal{R}^{\mathsf{t}}\ w\,x \quad x\ \text{fresh}}{v\ \mathcal{R}^{\mathsf{v}}\ w} \qquad \frac{E[x]\ \mathcal{R}^{\mathsf{t}}\ F[x] \quad x\ \text{fresh}}{E\ \mathcal{R}^{\mathsf{c}}\ F} \qquad \frac{\langle h \mid E[x] \rangle \in \mathcal{R}{\uparrow} \quad x, h\ \text{fresh}}{E \in \mathcal{R}{\uparrow}^{\mathsf{c}}}$$

We use these extensions to define progress as follows.

Definition 2. *A candidate relation $\mathcal{R}$ progresses to $\mathcal{S}$, $\mathcal{T}$ written $\mathcal{R} \rightarrowtail \mathcal{S}, \mathcal{T}$, if $\mathcal{R} \subseteq \mathcal{S}$, $\mathcal{S} \subseteq \mathcal{T}$, and*

1. *$c\ \mathcal{R}\ d$ implies*
   - *if $c \to c'$, then $d \to^* d'$ and $c'\ \mathcal{T}\ d'$;*
   - *if $c = \langle h \mid v \rangle$, then $d \to^* \langle g \mid w \rangle$, $h\ \mathcal{S}^{\mathsf{h}}\ g$, and $v\ \mathcal{S}^{\mathsf{v}}\ w$;*
   - *if $c = \langle h \mid E[x\,v] \rangle$, then either*
     - *$d \to^* \langle g \mid F[x\,w] \rangle$, $h\ \mathcal{T}^{\mathsf{h}}\ g$, $E\ \mathcal{T}^{\mathsf{c}}\ F$, and $v\ \mathcal{T}^{\mathsf{v}}\ w$, or*
     - *$E \in \mathcal{T}{\uparrow}^{\mathsf{c}}$.*
2. *$c \in \mathcal{R}{\uparrow}$ implies $c \neq \langle h \mid v \rangle$ for all h and v and*
   - *if $c \to c'$, then $c' \in \mathcal{T}{\uparrow}$;*
   - *if $c = \langle h \mid E[x\,v] \rangle$, then $E \in \mathcal{T}{\uparrow}^{\mathsf{c}}$.*

A normal-form simulation is a candidate relation $\mathcal{R}$ such that $\mathcal{R} \rightarrowtail \mathcal{R}, \mathcal{R}$, and a bisimulation is a candidate relation $\mathcal{R}$ such that $\mathcal{R}$ and $\mathcal{R}^{-1}$ are simulations. Normal-form bisimilarity $\approx$ is the union of all normal-form bisimulations.

We test values and contexts by applying or plugging them with a fresh variable x, and running them in a fresh store; with a global memory, the value represented by x may access any reference and assign it an arbitrary value, hence the need for a fresh store. The stores of two bisimilar value configurations must have the same domain, as it would be easy to distinguish them otherwise by testing the content of the references that would be in one store but not in the other.

The main novelty compared to usual definitions of normal-form bisimilarity [3,11] is the set of (deferred) diverging configurations used in the stuck terms

clause. We detect that E in a configuration $\langle h \mid E[x\,v]\rangle$ is (deferred) diverging by running $\langle h' \mid E[y]\rangle$ where y and h' are fresh; this configuration may then diverge or evaluate to an other deferred diverging configuration $\langle h \mid E'[x\,v]\rangle$.

Like in the plain λ-calculus [3], $\mathcal{R}$ progresses towards $\mathcal{S}$ in the value clause and $\mathcal{T}$ in the others; the former is passive while the others are active. Our framework prevents some up-to techniques from being applied after a passive transition. In particular, we want to forbid the application of bisimulation up to context as it would be unsound: we could deduce that $v\,x$ and $w\,x$ are equivalent for all v and w just by building a candidate relation containing v and w.

Example 1. To prove that $\langle h \mid x\,v\,\Omega\rangle \approx \langle h \mid \Omega\rangle$ holds for all v and h, we prove that $\mathcal{R} \stackrel{\mathsf{def}}{=} \{(\langle h \mid x\,v\,\Omega\rangle, \langle h \mid \Omega\rangle)), \{\langle g \mid y\,\Omega\rangle \mid y, g \text{ fresh}\}\}$ is a bisimulation. Indeed, $\langle h \mid x\,v\,\Omega\rangle$ is stuck with $\langle g \mid y\,\Omega\rangle \in \mathcal{R}\!\!\uparrow$ for fresh y and g, and we have $\langle g \mid y\,\Omega\rangle \to \langle g \mid y\,\Omega\rangle$. Conversely, the transition $\langle h \mid \Omega\rangle \to \langle h \mid \Omega\rangle$ is matched by $\langle h \mid x\,v\,\Omega\rangle \to^* \langle h \mid x\,v\,\Omega\rangle$ and the resulting terms are in $\mathcal{R}$.

2.3   Soundness

In this framework, proving that $\approx$ is sound is a consequence that a form of bisimulation up to context is valid, a result which itself may require to prove that other up-to techniques are valid. We distinguish the techniques which can be used in passive clauses (called *strong* up-to techniques), from the ones which cannot. An up-to technique (resp. strong up-to technique) is a function f such that $\mathcal{R} \rightarrowtail \mathcal{R}, f(\mathcal{R})$ (resp. $\mathcal{R} \rightarrowtail f(\mathcal{R}), f(\mathcal{R})$) implies $\mathcal{R} \subseteq \approx$. To show that a given f is an up-to technique, we rely on a notion of *respectfulness*, which is simpler to prove and gives sufficient conditions for f to be an up-to technique.

We briefly recall the notions we need from our previous work [2]. We extend $\subseteq$ and $\cup$ to functions argument-wise (e.g., $(f \cup g)(\mathcal{R}) = f(\mathcal{R}) \cup g(\mathcal{R})$), and given a set $\mathfrak{F}$ of functions, we also write $\mathfrak{F}$ for the function defined as $\bigcup_{f \in \mathfrak{F}} f$. We define f^ω as $\bigcup_{n \in \mathbb{N}} f^n$. We write id for the identity function on relations, and $\widehat{f}$ for $f \cup \mathsf{id}$. A function f is monotone if $\mathcal{R} \subseteq \mathcal{S}$ implies $f(\mathcal{R}) \subseteq f(\mathcal{S})$. We write $\mathcal{P}_{\mathit{fin}}(\mathcal{R})$ for the set of finite subsets of $\mathcal{R}$, and we say f is continuous if it can be defined by its image on these finite subsets, i.e., if $f(\mathcal{R}) \subseteq \bigcup_{\mathcal{S} \in \mathcal{P}_{\mathit{fin}}(\mathcal{R})} f(\mathcal{S})$. The up-to techniques we use are defined by inference rules with a finite number of premises, so they are trivially continuous.

Definition 3. *A function f evolves to g, h, written $f \rightsquigarrow g, h$, if for all $\mathcal{R}$ and $\mathcal{T}$, $\mathcal{R} \rightarrowtail \mathcal{R}, \mathcal{T}$ implies $f(\mathcal{R}) \rightarrowtail g(\mathcal{R}), h(\mathcal{T})$. A function f strongly evolves to g, h, written $f \rightsquigarrow_{\mathsf{s}} g, h$, if for all $\mathcal{R}$, $\mathcal{S}$, and $\mathcal{T}$, $\mathcal{R} \rightarrowtail \mathcal{S}, \mathcal{T}$ implies $f(\mathcal{R}) \rightarrowtail g(\mathcal{S}), h(\mathcal{T})$.*

Evolution can be seen as progress for functions on relations. Evolution is more restrictive than strong evolution, as it requires $\mathcal{R}$ such that $\mathcal{R} \rightarrowtail \mathcal{R}, \mathcal{T}$.

Definition 4. *A set $\mathfrak{F}$ of continuous functions is* respectful *if there exists $\mathfrak{S}$ such that $\mathfrak{S} \subseteq \mathfrak{F}$ and*

- *for all $f \in \mathfrak{S}$, we have $f \rightsquigarrow_{\mathsf{s}} \widehat{\mathfrak{S}}^\omega, \widehat{\mathfrak{F}}^\omega$;*
- *for all $f \in \mathfrak{F}$, we have $f \rightsquigarrow \widehat{\mathfrak{S}}^\omega \circ \widehat{\mathfrak{F}} \circ \widehat{\mathfrak{S}}^\omega, \widehat{\mathfrak{F}}^\omega$.*

$$\frac{c \; \mathcal{R} \; d \qquad v \; \mathcal{R}^{\mathsf{v}} \; w}{c\{v/x\} \; \mathsf{subst}(\mathcal{R}) \; d\{w/x\}} \qquad \frac{c \in \mathcal{R}{\uparrow}}{c\{v/x\} \in \mathsf{subst}(\mathcal{R}){\uparrow}} \qquad \frac{\langle h \mid t \rangle \; \mathcal{R} \; \langle g \mid s \rangle \qquad E \; \mathcal{R}^{\mathsf{c}} \; F}{\langle h \mid E[t] \rangle \; \mathsf{plug}_{\mathsf{c}}(\mathcal{R}) \; \langle g \mid F[s] \rangle}$$

$$\frac{\langle h \mid t \rangle \in \mathcal{R}{\uparrow}}{\langle h \mid E[t] \rangle \in \mathsf{plug}_{\uparrow}(\mathcal{R}){\uparrow}} \qquad \frac{c \to^* c' \qquad d \to^* d' \qquad c' \; \mathcal{R} \; d'}{c \; \mathsf{red}(\mathcal{R}) \; d}$$

$$\frac{c \in \mathcal{R}{\uparrow}}{c \; \mathsf{div}(\mathcal{R}) \; d} \qquad \frac{E \in \mathcal{R}{\uparrow}^{\mathsf{c}}}{\langle h \mid E[t] \rangle \in \mathsf{plugdiv}(\mathcal{R}){\uparrow}}$$

Fig. 1. Up-to techniques for the calculus with global store

In words, a function is in a respectful set $\mathfrak{F}$ if it evolves towards a combination of functions in $\mathfrak{F}$ after active clauses, and in $\mathfrak{S}$ after passive ones. When checking that f is regular (second case), we can use a regular function at most once after a passive clause. The (possibly empty) subset $\mathfrak{S}$ intuitively represents the strong up-to techniques of $\mathfrak{F}$. If $\mathfrak{S}_1$ and $\mathfrak{S}_2$ are subsets of $\mathfrak{F}$ which verify the conditions of the definition, then $\mathfrak{S}_1 \cup \mathfrak{S}_2$ also does, so there exists the largest subset of $\mathfrak{F}$ which satisfies the conditions, written $\mathsf{strong}(\mathfrak{F})$.

Lemma 1. *Let $\mathfrak{F}$ be a respectful set.*

- *If $f \in \mathfrak{F}$, then f is an up-to technique. If $f \in \mathsf{strong}(\mathfrak{F})$, then f is a strong up-to technique.*
- *For all $f \in \mathfrak{F}$, we have $f(\approx) \subseteq \approx$.*

Showing that f is in a respectful set $\mathfrak{F}$ is easier than proving it is an up-to technique. Besides, proving that a bisimulation up to context is respectful implies that $\approx$ is preserved by contexts thanks to the last property of Lemma 1.

The up-to techniques for the calculus with global store are given in Fig. 1. The techniques subst and plug allow to prove that $\approx$ is preserved by substitution and by evaluation contexts. The remaining ones are auxiliary techniques which are used in the respectfulness proof: red relies on the fact that the calculus is deterministic to relate terms up to reduction steps. The technique div allows to relate a diverging configuration to any other configuration, while $\mathsf{plugdiv}$ states that if E is a diverging context, then $\langle h \mid E[t] \rangle$ is a diverging configuration for all h and t. We distinguish the technique $\mathsf{plug}_{\mathsf{c}}$ from $\mathsf{plug}_{\uparrow}$ to get a more fine-grained classification, as $\mathsf{plug}_{\mathsf{c}}$ is the only one which is not strong.

Lemma 2. *The set $\mathfrak{F} \stackrel{\mathsf{def}}{=} \{\mathsf{subst}, \mathsf{plug}_m, \mathsf{red}, \mathsf{div}, \mathsf{plugdiv} \mid m \in \{\mathsf{c}, \uparrow\}\}$ is respectful, with $\mathsf{strong}(\mathfrak{F}) = \mathfrak{F} \setminus \{\mathsf{plug}_{\mathsf{c}}\}$.*

We omit the proof, as it is similar but much simpler than for the calculus with local store of Sect. 3. We deduce that $\approx$ is sound using Lemma 1.

Theorem 1. *For all t, s, and fresh store h, if $\langle h \mid t \rangle \approx \langle h \mid s \rangle$, then $t \equiv s$.*

2.4   Completeness

We prove the reverse implication by building a bisimulation which contains $\equiv$.

Theorem 2. *For all t, s, if $t \equiv s$, then for all fresh stores h, $\langle h \mid t \rangle \approx \langle h \mid s \rangle$.*

Proof (Sketch). It suffices to show that the candidate $\mathcal{R}$ defined as

$$\{(\langle h \mid t \rangle, \langle g \mid s \rangle) \mid \forall E, h_E, \text{ closing } \int, \langle h \uplus h_E \mid E[t] \rangle \int \Downarrow \ \Rightarrow \ \langle g \uplus h_E \mid E[s] \rangle \int \Downarrow \}$$
$$\cup \ \{\langle h \mid t \rangle \mid \forall E, h_E, \text{ closing } \int, \langle h \uplus h_E \mid E[t] \rangle \int \Uparrow \}$$

is a simulation. We proceed by case analysis on the behavior of $\langle h \mid t \rangle$. The details are in the report [4]; we sketch the proof in the case when $\langle h \mid t \rangle \ \mathcal{R} \ \langle g \mid s \rangle$, $t = E[x \, v]$, and E is not deferred diverging.

A first step is to show that $\langle g \mid s \rangle$ also evaluates to an open-stuck configuration with x in function position. To do so, we consider a fresh l and we define $\int$ such that $\int(y)$ sets l at 1 when it is first applied if $y = x$, and at 2 if $y \neq x$. Then $\langle h \uplus l := 0 \mid t \rangle \int$ sets l at 1, which should also be the case of $\langle g \uplus l := 0 \mid s \rangle \int$, and it is possible only if $\langle g \mid s \rangle \to^* \langle g' \mid F[x \, w] \rangle$ for some g', F, and w.

We then have to show that $E \ \mathcal{R}^c \ F$, $v \ \mathcal{R}^v \ w$, and $h \ \mathcal{R}^h \ g'$. We sketch the proof for the contexts, as the proofs for the values and the stores are similar. Given h_f a fresh store, y a fresh variable, E' a context, $h_{E'}$ a store, $\int$ a closing substitution, we want $\langle h_f \uplus h_{E'} \mid E'[E[y]] \rangle \int \Downarrow$ iff $\langle h_f \uplus h_{E'} \mid E'[F[y]] \rangle \int \Downarrow$.

Let l be a fresh reference. Assuming $\mathsf{dom}(h) = \{l_1 \ldots l_n\}$, given a term t, we write $\bigcup_i l_i := h; t$ for $l_1 := h(l_1); \ldots l_n := h(l_n); t$. We define

$$\int_x \ \overset{\text{def}}{=} \ \begin{cases} x \mapsto \lambda a.\text{if } !l = 0 \text{ then } l := 1; \bigcup_i l_i := h_f \uplus h_{E'}; \int(y) \text{ else } \int(x) \, a \\ z \mapsto \int'(z) \quad \text{if } z \neq x \end{cases}$$

The substitution $\int_x$ behaves like $\int$ except that when $\int_x(x)$ is applied for the first time, it replaces its argument by $\int(y)$ and sets the store to $h_f \uplus h_{E'}$. Therefore $\langle h \uplus l := 0 \mid E'[t] \rangle \int_x \ \to^* \ \langle h_f \uplus h_{E'} \uplus l := 1 \mid E'[E[y]] \rangle \int_x$, but this configuration then behaves like $\langle h_f \uplus h_{E'} \mid E'[E[y]] \rangle \int$. Similarly, $\langle g \uplus l := 0 \mid E'[s] \rangle \int_x$ evaluates to a configuration equivalent to $\langle h_f \uplus h_{E'} \mid E'[F[y]] \rangle \int$, and since $\langle h \uplus l := 0 \mid E'[t] \rangle \int_x \Downarrow$ implies $\langle g \uplus l := 0 \mid E'[s] \rangle \int_x \Downarrow$, we can conclude from there.

3   Local Store

We adapt the ideas of the previous section to a calculus where terms create their own local store. To be able to deal with local resources, the relation we define mixes principles from normal-form and environmental bisimilarities.

3.1   Syntax, Semantics, and Contextual Equivalence

In this section, the terms no longer share a global store, but instead must create local references before storing values. We extend the syntax of Sect. 2 with a construct to create a new reference.

Terms: $t, s ::= \dots \mid$ new $l := v$ in t

Reference creation new $l := v$ in t binds l in t; we identify terms up to α-conversion of their references. We write $\mathsf{fr}(t)$ and $\mathsf{fr}(E)$ for the set of free references of t or E, and a term or context is *reference-closed* if its set of free references is empty. Following [18] and in contrast with [5,6], references are not values, but we can still give access to a reference l by passing $\lambda x.!l$ and $\lambda x.l := x; \lambda y.y$.

As before, the semantics is defined on configurations $\langle h \mid t \rangle$ verifying $\mathsf{fr}(t) \subseteq \mathsf{dom}(h)$ and for all $l \in \mathsf{dom}(h)$, $\mathsf{fr}(h(l)) \subseteq \mathsf{dom}(h)$. We add to the rules of Sect. 2 the following one for reference creation.

$$\langle h \mid \text{new } l := v \text{ in } t \rangle \rightarrow \langle h \uplus l := v \mid t \rangle$$

We remind that $\uplus$ is defined for disjoint stores only, so the above rule assumes that $l \notin \mathsf{dom}(h)$, which is always possible using α-conversion.

We define contextual equivalence on reference-closed terms as we expect programs to allocate their own store.

Definition 5. *Two reference-closed terms t and s are contextually equivalent, written $t \equiv s$, if for all reference-closed evaluation contexts E and closing substitutions $\int$, $\langle \emptyset \mid E[t] \rangle \int \Downarrow$ iff $\langle \emptyset \mid E[s] \rangle \int \Downarrow$.*

3.2   Bisimilarity

With local stores, an external observer no longer has direct access to the stored values. In presence of such information hiding, a sound bisimilarity relies on an *environment* to accumulate terms which should be tested in different stores [8].

Example 2. Let $f_1 \stackrel{\text{def}}{=} \lambda x.$if $!l =$ true then $l :=$ false; true else false and $f_2 \stackrel{\text{def}}{=} \lambda x.$true. If we compare new $l :=$ true in f_1 and f_2 only once in the empty store, they would be seen as equivalent as they both return true, however f_1 modify its store, so running f_1 and f_2 a second time distinguishes them.

Environments generally contain only values [17], except in $\lambda\mu\rho$ [18], where plugged evaluation contexts are kept in the environment when comparing open-stuck configurations. In contrast with $\lambda\mu\rho$, our environment collects values, and we use a *stack* for registering contexts [7,10]. Unlike values, contexts are therefore tested only once, following a last-in first-out ordering. The next example shows that considering contexts repeatedly would lead to an overly-discriminating bisimilarity. For the stack discipline of testing contexts in action see Example 8 in Sect. 3.4.

Example 3. With the same f_1 and f_2 as in Example 2, the terms $t \stackrel{\text{def}}{=} \text{new } l :=$ true in $f_1 (x \, \lambda y.y)$ and $s \stackrel{\text{def}}{=} f_2 (x \, \lambda y.y)$ are contextually equivalent. Roughly, for all closing substitution $\int$, t and s either both diverge (if $\int(x) \, \lambda y.y$ diverges), or evaluate to true, since $\int(x)$ cannot modify the value in l. Testing $f_1 \, \square$ and $f_2 \, \square$ twice would discriminate them and wrongfully distinguish t and s.

Remark 1. The bisimilarity for $\lambda \mu \rho$ runs evaluation contexts several times and is still complete because of the μ operator, which, like call/cc, captures evaluation contexts, and may then execute them several times.

We let $\mathcal{E}$ range over sets of pairs of values, and ϵ over sets of values. Similarly, we write Σ for a stack of pairs of evaluation contexts and σ for a stack of evaluation contexts. We write $\odot$ for the empty stack, $::$ for the operator putting an element on top of a stack, and $+\!\!+$ for the concatenation of two stacks. The projection operator π_1 transforms a set or stack of pairs into respectively a set or stack of single elements by taking the first element of each pair. A candidate relation $\mathcal{R}$ can be composed of:

- quadruples $(\mathcal{E}, \Sigma, c, d)$, written $\mathcal{E}, \Sigma \vdash c \, \mathcal{R} \, d$, meaning that c and d are related under $\mathcal{E}$ and Σ;
- quadruples $(\mathcal{E}, \Sigma, h, g)$, written $\mathcal{E}, \Sigma \vdash h \, \mathcal{R} \, g$, meaning that the elements of $\mathcal{E}$ and the top of Σ should be related when run with the stores h and g;
- triples (ϵ, σ, c), written $\epsilon, \sigma \vdash c \in \mathcal{R}\!\!\uparrow$, meaning that either c is (deferred) diverging, or σ is non-empty and contains a (deferred) diverging context;
- triples (ϵ, σ, h), written $\epsilon, \sigma \vdash h \in \mathcal{R}\!\!\uparrow$, meaning that σ is non-empty and contains a (deferred) diverging context.

Definition 6. *A candidate relation $\mathcal{R}$ progresses to $\mathcal{S}, \mathcal{T}$ written $\mathcal{R} \rightarrowtail \mathcal{S}, \mathcal{T}$, if $\mathcal{R} \subseteq \mathcal{S}, \mathcal{S} \subseteq \mathcal{T}$, and*

1. *$\mathcal{E}, \Sigma \vdash c \, \mathcal{R} \, d$ implies*
    - *if $c \rightarrow c'$, then $d \rightarrow^* d'$ and $\mathcal{E}, \Sigma \vdash c' \, \mathcal{T} \, d'$;*
    - *if $c = \langle h \mid v \rangle$, then either*
        - *$d \rightarrow^* \langle g \mid w \rangle$, and $\mathcal{E} \cup \{(v, w)\}, \Sigma \vdash h \, \mathcal{S} \, g$, or*
        - *$\Sigma \neq \odot$ and $\pi_1(\mathcal{E}) \cup \{v\}, \pi_1(\Sigma) \vdash h \in \mathcal{S}\!\!\uparrow$;*
    - *if $c = \langle h \mid E[x \, v] \rangle$, then either*
        - *$d \rightarrow^* \langle g \mid F[x \, w] \rangle$, and $\mathcal{E} \cup \{(v, w)\}, (E, F) :: \Sigma \vdash h \, \mathcal{S} \, g$, or*
        - *$\pi_1(\mathcal{E}) \cup \{v\}, E :: \pi_1(\Sigma) \vdash h \in \mathcal{S}\!\!\uparrow$.*
2. *$\mathcal{E}, \Sigma \vdash h \, \mathcal{R} \, g$ implies*
    - *if $v \, \mathcal{E} \, w$, then $\mathcal{E}, \Sigma \vdash \langle h \mid v \, x \rangle \, \mathcal{S} \, \langle g \mid w \, x \rangle$ for a fresh x;*
    - *if $\Sigma = (E, F) :: \Sigma'$, then $\mathcal{E}, \Sigma' \vdash \langle h \mid E[x] \rangle \, \mathcal{S} \, \langle g \mid F[x] \rangle$ for a fresh x.*
3. *$\epsilon, \sigma \vdash c \in \mathcal{R}\!\!\uparrow$ implies*
    - *if $c \rightarrow c'$, then $\epsilon, \sigma \vdash c' \in \mathcal{T}\!\!\uparrow$;*
    - *if $c = \langle h \mid v \rangle$, then $\sigma \neq \odot$ and $\epsilon \cup \{v\}, \sigma \vdash h \in \mathcal{S}\!\!\uparrow$;*
    - *if $c = \langle h \mid E[x \, v] \rangle$, then $\epsilon \cup \{v\}, E :: \sigma \vdash h \in \mathcal{S}\!\!\uparrow$.*
4. *$\epsilon, \sigma \vdash h \in \mathcal{R}\!\!\uparrow$ implies that $\sigma \neq \odot$ and*
    - *if $v \in \epsilon$, then $\epsilon, \sigma \vdash \langle h \mid v \, x \rangle \in \mathcal{S}\!\!\uparrow$ for a fresh x;*
    - *if $\sigma = E :: \sigma'$, then $\epsilon, \sigma' \vdash \langle h \mid E[x] \rangle \in \mathcal{S}\!\!\uparrow$ for a fresh x.*

A normal-form simulation is a candidate relation $\mathcal{R}$ such that $\mathcal{R} \rightarrowtail \mathcal{R}, \mathcal{R}$, and a bisimulation is a candidate relation $\mathcal{R}$ such that $\mathcal{R}$ and $\mathcal{R}^{-1}$ are simulations. Normal-form bisimilarity $\approx$ is the union of all normal-form bisimulations.

When $\mathcal{E}, \Sigma \vdash c \mathcal{R} d$, we reduce c until we get a value v or a stuck term $E[x\,v]$. At that point, either d also reduces to a normal form of the same kind, or we test (the first projection of) the stack Σ for divergence, assuming it is not empty. In the former case, we add the values to $\mathcal{E}$ and the evaluation contexts at the top of Σ, getting a judgment of the form $\mathcal{E}', \Sigma' \vdash h \mathcal{R} g$, which then tests the environment and the stack by running either terms in $\mathcal{E}'$ or at the top of Σ'.

Example 4. We sketch the bisimulation proof for the terms t and s of Example 3. Because $\langle \emptyset \mid t \rangle \rightarrow^* \langle l := \mathsf{true} \mid f_1\,(x\,\lambda y.y) \rangle$ and $\langle \emptyset \mid s \rangle = \langle \emptyset \mid f_2\,(x\,\lambda y.y) \rangle$, we need to define $\mathcal{R}$ such that $\{(\lambda y.y, \lambda y.y)\}, (f_1\,\square, f_2\,\square) :: \odot \vdash l := \mathsf{true} \mathcal{R} \emptyset$. Testing the equal values in the environment is easy with up-to techniques. For the contexts on the stack, we need $\{(\lambda y.y, \lambda y.y)\}, \odot \vdash \langle l := \mathsf{true} \mid f_1\,z \rangle \mathcal{R} \langle \emptyset \mid f_2\,z \rangle$ for a fresh z. Since $\langle l := \mathsf{true} \mid f_1\,z \rangle \rightarrow^* \langle l := \mathsf{false} \mid \mathsf{true} \rangle$ and $\langle \emptyset \mid f_2\,z \rangle \rightarrow^* \langle \emptyset \mid \mathsf{true} \rangle$, we need $\{(\lambda y.y, \lambda y.y), (\mathsf{true}, \mathsf{true})\}, \odot \vdash l := \mathsf{false} \mathcal{R} \emptyset$, which is simple to check.

Example 5. In contrast, we show that $t' \stackrel{\mathsf{def}}{=} \mathsf{new}\ l := \mathsf{true}\ \mathsf{in}\ f_1\,(x\,\lambda y.l := y; y)$ and $s' \stackrel{\mathsf{def}}{=} f_2\,(x\,\lambda y.y)$ are not bisimilar. We would need to build $\mathcal{R}$ such that $\{(\lambda y.l := y; y, \lambda y.y)\}, (f_1\,\square, f_2\,\square) :: \odot \vdash l := \mathsf{true} \mathcal{R} \emptyset$. Testing the values in the environment, we want $\{(\lambda y.l := y; y, \lambda y.y), (z, z)\}, (f_1\,\square, f_2\,\square) :: \odot \vdash l := z \mathcal{R} \emptyset$ for a fresh z. Executing the contexts on the stack, we get a stuck term of the form $\mathsf{if}\ z\ \mathsf{then}\ l := \mathsf{false}; \mathsf{true}\ \mathsf{else}\ \mathsf{false}$ and a value true, which cannot be related, because the former is not deferred diverging.

The terms t' and s' are therefore not bisimilar, and they are indeed not contextually equivalent, since t' gives access to its private reference by passing $\lambda y.l := y; y$ to x. The function represented by x can then change the value of l to false and break the equivalence.

The last two cases of the bisimulation definition aim at detecting a deferred diverging context. The judgment $\epsilon, \sigma \vdash h \in \mathcal{R}{\uparrow}$ roughly means that if $\sigma = E_n :: \ldots E_1 :: \odot$, then the configuration $\langle h' \mid E_1[\ldots E_n[x]] \rangle$ diverges for all fresh x and all h' obtained by running a term from $\mathcal{E}$ with the store h. As a result, when $\epsilon, \sigma \vdash h \in \mathcal{R}{\uparrow}$, we have two possibilities: either we run a term from $\mathcal{E}$ in h to potentially change h, or we run the context at the top of σ (which cannot be empty in that case) to check if it is diverging. In both cases, we get a judgment of the form $\epsilon, \sigma' \vdash c \in \mathcal{R}{\uparrow}$. In that case, either c diverges and we are done, or it terminates, meaning that we have to look for divergence in σ'.

Example 6. We prove that $\langle \emptyset \mid x\,v\,\Omega \rangle$ and $\langle \emptyset \mid \Omega \rangle$ are bisimilar. We define $\mathcal{R}$ such that $\emptyset, \odot \vdash \langle \emptyset \mid x\,v\,\Omega \rangle \mathcal{R} \langle \emptyset \mid \Omega \rangle$, for which we need $\{v\}, \square\,\Omega :: \odot \vdash \emptyset \in \mathcal{R}{\uparrow}$, which itself holds if $\{v\}, \odot \vdash \langle \emptyset \mid y\,\Omega \rangle \in \mathcal{R}{\uparrow}$.

Finally, only the two clauses where a reduction step takes place are active; all the others are passive, because they are simply switching from one judgment to

$$\frac{\mathcal{E}, \Sigma \vdash c \; \mathcal{R} \; d \qquad v \; \mathcal{E} \; w \qquad x \notin \mathsf{fv}(v) \cup \mathsf{fv}(w)}{\mathcal{E}\{(v,w)/x\}, \Sigma\{(v,w)/x\} \vdash c\{v/x\} \; \mathsf{subst_c}(\mathcal{R}) \; d\{w/x\}}$$

$$\frac{\mathcal{E}, \Sigma_1 \mathbin{+\!\!+} (E_1, F_1) :: (E_2, F_2) :: \Sigma_2 \vdash \langle h \mid t \rangle \; \mathcal{R} \; \langle g \mid s \rangle}{\mathcal{E}, \Sigma_1 \mathbin{+\!\!+} (E_2[E_1], F_2[F_1]) :: \Sigma_2 \vdash \langle h \mid t \rangle \; \mathsf{ccomp}(\mathcal{R}) \; \langle g \mid s \rangle}$$

$$\frac{\mathcal{E}, (E, F) :: \Sigma \vdash \langle h \mid t \rangle \; \mathcal{R} \; \langle g \mid s \rangle}{\mathcal{E}, \Sigma \vdash \langle h \mid E[t] \rangle \; \mathsf{plug}(\mathcal{R}) \; \langle g \mid F[s] \rangle} \qquad \frac{c \to^* c' \qquad d \to^* d' \qquad \mathcal{E}, \Sigma \vdash c' \; \mathcal{R} \; d'}{\mathcal{E}, \Sigma \vdash c \; \mathsf{red}(\mathcal{R}) \; d}$$

$$\frac{\epsilon, \sigma \vdash \langle h \mid t \rangle \in \mathcal{R}{\uparrow} \qquad \pi_1(\mathcal{E}) = \epsilon \qquad \pi_1(\Sigma) = \sigma}{\mathcal{E}, \Sigma \vdash \langle h \mid t \rangle \; \mathsf{div}(\mathcal{R}) \; \langle g \mid s \rangle} \qquad \frac{\mathcal{E}, \Sigma \vdash c \; \mathcal{R} \; d \qquad \mathcal{E}' \subseteq \mathcal{E}}{\mathcal{E}', \Sigma \vdash c \; \mathsf{weak}(\mathcal{R}) \; d}$$

$$\frac{\mathcal{E}, \Sigma_1 \mathbin{+\!\!+} \Sigma_2 \vdash \langle h \mid t \rangle \; \mathcal{R} \; \langle g \mid s \rangle \qquad \mathsf{fr}(E) \subseteq \mathsf{dom}(h')}{\mathcal{E}, \Sigma_1 \mathbin{+\!\!+} (E, E) :: \Sigma_2 \vdash \langle h \uplus h' \mid t \rangle \; \mathsf{refl}(\mathcal{R}) \; \langle g \uplus h' \mid s \rangle}$$

Fig. 2. Selected up-to techniques for the calculus with local store

the other without any real progress taking place. For example, when comparing value configurations, we go from a configuration judgment $\mathcal{E}, \Sigma \vdash c \; \mathcal{R} \; d$ to a store judgment $\mathcal{E}, \Sigma \vdash h \; \mathcal{R} \; g$ or a diverging store judgment $\mathcal{E}, \Sigma \vdash h \in \mathcal{R}{\uparrow}$. In a (diverging) store judgment, we simply decide whether we reduce a term from the store of from the stack, going back to a (diverging) configuration judgment. Actual progress is made only when we start reducing the chosen configuration.

3.3   Soundness and Completeness

We briefly discuss the up-to techniques we need to prove soundness. We write $\mathcal{E}\{(v,w)/x\}$ for the environment $\{(v'\{v/x\}, w'\{w/x\}) \mid v' \; \mathcal{E} \; w'\}$, and we also define $\Sigma\{(x,w)/x\}$, $\epsilon\{v/x\}$, and $\sigma\{v/x\}$ as expected. To save space, Fig. 2 presents the up-to techniques for the configuration judgment only; see the report [4] for the other judgments.

As in Sect. 2.3, the techniques subst and plug allow to reason up to substitution and plugging into an evaluation context, except that the substituted values and plugged contexts must be taken from respectively the environment and the top of the stack. The technique div relates a diverging configuration to any configuration, like in the calculus with global store. The technique ccomp allows to merge successive contexts in the stack into one. The weakening technique weak, originally known as bisimulation up to environment [17], is an usual technique for environmental bisimulations. Making the environment smaller creates a weaker judgment, as having less testing terms means a less discriminating candidate relation. Bisimulation up to reduction red is also standard and allows for a big-step reasoning by ignoring reduction steps. Finally, the technique refl allows to introduce identical contexts in the stack, but also values in the environment or terms in configurations (see the report [4]).

We denote by $\mathsf{subst_c}$ the up to substitution technique restricted to the configuration and diverging configuration judgments, and by $\mathsf{subst_s}$ the restriction to the store and diverging store judgments.

Lemma 3. *The set $\mathfrak{F} \stackrel{\mathsf{def}}{=} \{\mathsf{subst}_m, \mathsf{plug}, \mathsf{ccomp}, \mathsf{div}, \mathsf{weak}, \mathsf{red}, \mathsf{refl} \mid m \in \{\mathsf{c}, \mathsf{s}\}\}$ is respectful, with $\mathsf{strong}(\mathfrak{F}) = \{\mathsf{subst_s}, \mathsf{ccomp}, \mathsf{div}, \mathsf{weak}, \mathsf{red}, \mathsf{refl}\}$.*

In contrast with Sect. 2.3 and our previous work [3], $\mathsf{subst_c}$ is *not* strong, because values are taken from the environment. Indeed, with $\mathsf{subst_c}$ strong, from $\{(v, w)\}, \odot \vdash \emptyset \; \mathcal{R} \; \emptyset$, we could derive $\{(v, w)\}, \odot \vdash \langle \emptyset \mid x \, y \rangle \; \mathsf{refl}(\mathcal{R}) \; \langle \emptyset \mid x \, y \rangle$ and then $\{(v, w)\}, \odot \vdash \langle \emptyset \mid v \, x \rangle \; \mathsf{subst_c}(\mathsf{refl}(\mathcal{R})) \; \langle \emptyset \mid w \, x \rangle$ for any v and w, which would be unsound.

The respectfulness proofs are in the report [4]. Using refl, plug, $\mathsf{subst_c}$, and Lemma 1 we prove that $\approx$ is preserved by evaluation contexts and substitution, from which we deduce it is sound w.r.t. contextual equivalence.

Theorem 3. *For all t and s, if $\emptyset, \odot \vdash \langle \emptyset \mid t \rangle \approx \langle \emptyset \mid s \rangle$, then $t \equiv s$.*

To establish completeness, we follow the proof of Theorem 2, i.e., we construct a candidate relation $\mathcal{R}$ that contains $\equiv$ and prove it is a simulation by case analysis on the behavior of the related terms.

Theorem 4. *For all t and s, if $t \equiv s$, then $\emptyset, \odot \vdash \langle \emptyset \mid t \rangle \approx \langle \emptyset \mid s \rangle$.*

The main difference is that the contexts and closing substitutions are built from the environment using compatible closures [17], to take into account the private resources of the related terms. We discuss the proof in the report [4].

3.4 Examples

Example 7. We start by the so-called awkward example [5,6,15]. Let

$$v \stackrel{\mathsf{def}}{=} \lambda f.l := 0; f \,(); l := 1; f \,(); !l \qquad w \stackrel{\mathsf{def}}{=} \lambda f.f \,(); f \,(); 1.$$

We equate new $l := 0$ in v and w, building the candidate $\mathcal{R}$ incrementally, starting from $\{(v, w)\}, \odot \vdash l := 0 \; \mathcal{R} \; \emptyset$.

Running v and w with a fresh variable f, we obtain $\langle l := 0 \mid E_1[f \,()] \rangle$ and $\langle \emptyset \mid E_2[f \,()] \rangle$ with $E_1 \stackrel{\mathsf{def}}{=} \square; l := 1; f \,(); !l$ and $F_1 \stackrel{\mathsf{def}}{=} \square; f \,(); 1$. Ignoring the identical unit arguments (using refl), we need $\{(v, w)\}, (E_1, F_1) :: \odot \vdash l := 0 \; \mathcal{R} \; \emptyset$; from that point, we can either test v and w again, resulting into an extra pair (E_1, F_1) on the stack, or run $\langle l := 0 \mid E_1[g] \rangle$ and $\langle \emptyset \mid F_1[g] \rangle$ for a fresh g instead.

In the latter case, we get $\langle l := 1 \mid E_2[g \,()] \rangle$ and $\langle \emptyset \mid F_2[g \,()] \rangle$, with $E_2 \stackrel{\mathsf{def}}{=} \square; !l$ and $F_2 \stackrel{\mathsf{def}}{=} \square; 1$, so we want $\{(v, w)\}, (E_2, F_2) :: \odot \vdash l := 1 \; \mathcal{R} \; \emptyset$ (ignoring again the units). From there, testing v and w produces $\{(v, w)\}, (E_1, F_1) :: (E_2, F_2) :: \odot \vdash l := 0 \; \mathcal{R} \; \emptyset$, while executing $\langle l := 1 \mid E_2[x] \rangle$ and $\langle \emptyset \mid F_2[x] \rangle$ for a fresh x gives us $\langle l := 1 \mid 1 \rangle$ and $\langle \emptyset \mid 1 \rangle$. This analysis suggests that $\mathcal{R}$ should be composed only of judgments of the form $\{(v, w)\}, \Sigma \vdash l := n \; \mathcal{R} \; \emptyset$ such that $n \in \{0, 1\}$ and

– Σ is an arbitrary stack composed only of pairs (E_1, F_1) or (E_2, F_2);
– if $\Sigma = (E_2, F_2) :: \Sigma'$, then $n = 1$.

We can check that such a candidate is a bisimulation, and it ensures that when l is read (when E_2 is executed), it contains the value 1.

Example 8. As a variation on the awkward example, let

$$v \overset{\text{def}}{=} \lambda f.l :=\ !l + 1; f\ (); l :=\ !l - 1; !l > 0 \qquad w \overset{\text{def}}{=} \lambda f.f\ (); \text{true}.$$

We show that $\langle \emptyset \mid \text{new } l := 1 \text{ in } v \rangle$ and $\langle \emptyset \mid w \rangle$ are bisimilar. Let $E \overset{\text{def}}{=} \square; l :=\ !l - 1; !l > 0$ and $F \overset{\text{def}}{=} \square; \text{true}$. We write $(E, F)^n$ for the stack $\odot$ if $n = 0$ and $(E, F) :: (E, F)^{n-1}$ otherwise. Then the candidate $\mathcal{R}$ verifying $\{(v, w)\}, (E, F)^n \vdash l := n + 1\ \mathcal{R}\ \emptyset$ for any n is a bisimulation. Indeed, running v and w increases the value stored in l and adds a pair (E, F) on the stack. If $n > 0$, we can run a copy of E and F, thus decreasing the value in l by 1, and then returning true in both cases.

Example 9. This deferred divergence example comes from Dreyer et al. [5]. Let

$$v_1 \overset{\text{def}}{=} \lambda x.\text{if } !l \text{ then } \Omega \text{ else } k := \text{true}; \lambda y.y \qquad w_1 \overset{\text{def}}{=} \lambda x.\Omega$$

$$v_2 \overset{\text{def}}{=} \lambda f.f\ v_1; \text{if } !k \text{ then } \Omega \text{ else } l := \text{true}; \lambda y.y \qquad w_2 \overset{\text{def}}{=} \lambda f.f\ w_1; \lambda y.y$$

We prove that $\text{new } l := \text{false in new } k := \text{false in } v_2$ is equivalent to w_2. Informally, if f in w_2 applies its argument w_1, the term diverges. Divergence also happens in v_2 but in a delayed fashion, as v_1 first sets k to true, and the continuation $t \overset{\text{def}}{=} \text{if } !k \text{ then } \Omega \text{ else } l := \text{true}; \lambda y.y$ then diverges. Similarly, if f stores w_1 or v_1 to later apply it, then divergence also occurs in both cases: in that case t sets l to true, and when v_1 is later applied, it diverges.

To build a candidate $\mathcal{R}$, we execute $\langle l := \text{false}; k := \text{false} \mid v_2\ f \rangle$ and $\langle \emptyset \mid w_2\ f \rangle$ for a fresh f, which gives us $\langle l := \text{false}; k := \text{false} \mid E[f\ v_1] \rangle$ and $\langle \emptyset \mid F[f\ w_1] \rangle$ with $E \overset{\text{def}}{=} \square; t$ and $F \overset{\text{def}}{=} \square; \lambda y.y$. We consider $\{(v_2, w_2), (v_1, w_1)\}, (E, F) :: \emptyset \vdash l := \text{false}; k := \text{false}\ \mathcal{R}\ \emptyset$, for which we have several checks to do. The interesting one is running $\langle l := \text{false}; k := \text{false} \mid v_1\ x \rangle$ and $\langle \emptyset \mid w_1\ x \rangle$, as we get $\langle l := \text{false}; k := \text{true} \mid \lambda y.y \rangle$ and $\langle \emptyset \mid \Omega \rangle$. In that case, we are showing that the stack contains divergence, by establishing that $\{v_2, v_1, \lambda y.y\}, E :: \emptyset \vdash l := \text{false}; k := \text{true} \in \mathcal{R}\uparrow$, and indeed, we have $\langle l := \text{false}; k := \text{true} \mid E[x] \rangle \rightarrow^* \langle l := \text{false}; k := \text{true} \mid \Omega \rangle$ for a fresh x. In the end, the relation $\mathcal{R}$ verifying

$$\{(v_2, w_2), (v_1, w_1)\}, (E, F)^n \vdash l := \text{false}; k := \text{false}\ \mathcal{R}\ \emptyset$$
$$\{(v_2, w_2), (v_1, w_1)\}, (E, F)^n \vdash \langle l := \text{false}; k := \text{true} \mid \lambda y.y \rangle\ \mathcal{R}\ \langle \emptyset \mid \Omega \rangle$$
$$\{v_2, v_1, \lambda y.y\}, E^n \vdash l := \text{false}; k := \text{true} \in \mathcal{R}\uparrow$$
$$\{v_2, v_1, \lambda y.y\}, E^n \vdash \langle l := \text{false}; k := \text{true} \mid \Omega \rangle \in \mathcal{R}\uparrow$$
$$\{(v_2, w_2), (v_1, w_1)\}, (E, F)^n \vdash l := \text{true}; k := \text{false}\ \mathcal{R}\ \emptyset$$
$$\{(v_2, w_2), (v_1, w_1)\}, (E, F)^n \vdash \langle l := \text{true}; k := \text{false} \mid \Omega \rangle\ \mathcal{R}\ \langle \emptyset \mid \Omega \rangle$$

for all n is a bisimulation up to refl and red.

4   Related Work and Conclusion

Related Work. As pointed out in Sect. 1, the other bisimilarities defined for state either feature universal quantification over testing arguments [9,12,17,19], or are complete only for a more expressive language [18]. Kripke logical relations [1,5] also involve quantification over arguments when testing terms of a functional type. Finally, denotational models [10,13] can also be used to prove program equivalence, by showing that the denotations of two terms are equal. However, computing such denotations is difficult in general, and the automation of this task is so far restricted to a language with first-order references [14].

The work most closely related to ours is Jaber and Tabareau's Kripke Open Bisimulation (KOB) [6]. A KOB tests functional terms with fresh variables and not with related values like a regular logical relation would do. To relate two given configurations, one has to provide a World Transition System (WTS) which states the invariants the heaps of the configurations should satisfy and how to go from one invariant to the other during the evaluation. Similarly, the bisimulations for the examples of Sect. 3.4 state properties which could be seen as invariants about the stores at different points of the evaluation.

The difficulty for KOB as well as with our bisimilarity is to come up with the right invariants about the heaps, expressed either as a WTS or as a bisimulation. We believe that choosing a technique over the other is just a matter of preference, depending on whether one is more comfortable with game semantics or with coinduction. It would be interesting to see if there is a formal correspondence between KOB and our bisimilarity; we leave this question as a future work.

Conclusion. We define a sound and complete normal-form bisimilarity for higher-order local state, with an environment to be able to run terms in different stores. We distinguish in the environment values which should be tested several times from the contexts which should be executed only once. The other difficulty is to relate deferred and regular diverging terms, which is taken care of by the specific judgments about divergence. The lack of quantification over arguments make the bisimulation proofs quite simple.

A future work would be to make these proofs even simpler by defining appropriate up-to techniques. The techniques we use in Sect. 3.3 to prove soundness turn out to be not that useful when establishing the equivalences of Sect. 3.4, except for trivial ones such as up to reduction or reflexivity. The difficulty in defining the candidate relations for the examples of Sect. 3.4 is in finding the right property relating the stack Σ to the store, so maybe an up-to technique could make this task easier.

As pointed out in Sect. 1, our results can be seen as an indication of what kind of additional infrastructure in a complete normal-form bisimilarity is required when the considered syntactic theory becomes less discriminative—in our case, when control operators vanish from the picture, and mutable state is the only extension of the λ-calculus. A question one could then ask is whether we can find a less expressive calculus—maybe the plain λ-calculus itself—for which a suitably enhanced normal-form bisimilarity is still complete.

References

1. Ahmed, A., Dreyer, D., Rossberg, A.: State-dependent representation independence. In: Pierce, B.C. (ed.) Proceedings of the Thirty-Fifth Annual ACM Symposium on Principles of Programming Languages, pp. 340–353. ACM Press, January 2009

2. Aristizábal, A., Biernacki, D., Lenglet, S., Polesiuk, P.: Environmental bisimulations for delimited-control operators with dynamic prompt generation. Logical Methods Comput. Sci. **13**(3) 2017

3. Biernacki, D., Lenglet, S., Polesiuk, P.: Proving soundness of extensional normal-form bisimilarities. In: Silva, A. (ed.) Proceedings of the 33rd Annual Conference on Mathematical Foundations of Programming Semantics (MFPS XXXIII), Ljubljana, Slovenia. Electronic Notes in Theoretical Computer Science, vol. 336, pp. 41–56, June 2017

4. Biernacki, D., Lenglet, S., Polesiuk, P.: A complete normal-form bisimilarity for state. Research report RR-9251, Inria, Nancy, France, January 2019

5. Dreyer, D., Neis, G., Birkedal, L.: The impact of higher-order state and control effects on local relational reasoning. J. Funct. Program. **22**(4–5), 477–528 (2012)

6. Jaber, G., Tabareau, N.: Kripke open bisimulation – a marriage of game semantics and operational techniques. In: Feng, X., Park, S. (eds.) APLAS 2015. LNCS, vol. 9458, pp. 271–291. Springer, Cham (2015)

7. Jagadeesan, R., Pitcher, C., Riely, J.: Open bisimulation for aspects. Trans. Aspect-Oriented Softw. Dev. **5**, 72–132 (2009)

8. Koutavas, V., Levy, P.B., Sumii, E.: From applicative to environmental bisimulation. In: Mislove, M., Ouaknine, J. (eds.) Proceedings of the 27th Annual Conference on Mathematical Foundations of Programming Semantics (MFPS XXVII), Pittsburgh, PA, USA. ENTCS, vol. 276, pp. 215–235, May 2011

9. Koutavas, V., Wand, M.: Small bisimulations for reasoning about higher-order imperative programs. In: Morrisett, J.G., Jones, S.L.P. (eds.) POPL 2006, Charleston, SC, USA, pp. 141–152. ACM Press (2006)

10. Laird, J.: A fully abstract trace semantics for general references. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 667–679. Springer, Heidelberg (2007)

11. Lassen, S.B.: Eager normal form bisimulation. In: Panangaden, P. (ed.) LICS 2005, Chicago, IL, pp. 345–354. IEEE Computer Society Press (2005)

12. Madiot, J.-M., Pous, D., Sangiorgi, D.: Bisimulations up-to: beyond first-order transition systems. In: Baldan, P., Gorla, D. (eds.) CONCUR 2014. LNCS, vol. 8704, pp. 93–108. Springer, Heidelberg (2014)

13. Murawski, A.S., Tzevelekos, N.: Game semantics for good general references. In: Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, pp. 75–84. IEEE Computer Society, June 2011

14. Murawski, A.S., Tzevelekos, N.: Algorithmic games for full ground references. Formal Methods Syst. Des. **52**(3), 277–314 (2018)

15. Pitts, A., Stark, I.: Operational reasoning for functions with local state. In: Gordon, A., Pitts, A. (eds.) Higher Order Operational Techniques in Semantics, pp. 227–273. Publications of the Newton Institute, Cambridge University Press (1998)

16. Sangiorgi, D.: The lazy lambda calculus in a concurrency scenario. In: Scedrov, A. (ed.) LICS 1992, Santa Cruz, California, pp. 102–109. IEEE Computer Society (1992)
17. Sangiorgi, D., Kobayashi, N., Sumii, E.: Environmental bisimulations for higher-order languages. ACM Trans. Program. Lang. Syst. **33**(1), 1–69 (2011)
18. Støvring, K., Lassen, S.B.: A complete, co-inductive syntactic theory of sequential control and state. In: Felleisen, M. (ed.) SIGPLAN Notices, POPL 2007, Nice, France, vol. 42, no. 1, pp. 161–172. ACM Press (2007)
19. Sumii, E.: A complete characterization of observational equivalence in polymorphic λ-calculus with general references. In: Grädel, E., Kahle, R. (eds.) CSL 2009. LNCS, vol. 5771, pp. 455–469. Springer, Heidelberg (2009)