Bisimulations up-to: beyond first-order transition systems

Jean-Marie Madiot, Damien Pous, Davide Sangiorgi

Bisimulations up-to: beyond first-order transition systems

Jean-Marie Madiot[1], Damien Pous[1], and Davide Sangiorgi[2]

[1] ENS Lyon, Université de Lyon, CNRS, INRIA, France,
[2] Università di Bologna, Italy, INRIA

Abstract. The bisimulation proof method can be enhanced by employing *'bisimulations up-to'* techniques. A comprehensive theory of such enhancements has been developed for first-order (i.e., CCS-like) labelled transition systems (LTSs) and bisimilarity, based on the notion of compatible function for fixed-point theory.
We transport this theory onto languages whose bisimilarity and LTS go beyond those of first-order models. The approach consists in exhibiting fully abstract translations of the more sophisticated LTSs and bisimilarities onto the first-order ones. This allows us to reuse directly the large corpus of up-to techniques that are available on first-order LTSs. The only ingredient that has to be manually supplied is the compatibility of basic up-to techniques that are specific to the new languages. We investigate the method on the π-calculus, the λ-calculus, and a (call-by-value) λ-calculus with references.

1 Introduction

One of the keys for the success of bisimulation is its associated proof method, whereby to prove two terms equivalent, one exhibits a relation containing the pair and one proves it to be a bisimulation. The bisimulation proof method can be enhanced by employing relations called *'bisimulations up-to'* [14, 18, 19]. These relations need not be bisimulations; they are simply *contained in* a bisimulation. Such techniques have been widely used in languages for mobility such as π-calculus or higher-order languages such as the λ-calculus, or Ambients (e.g., [22, 15, 11]).

Several forms of bisimulation enhancements have been introduced: 'bisimulation up-to bisimilarity' [16] where the derivatives obtained when playing bisimulation games can be rewritten using bisimilarity itself; 'bisimulation up-to transitivity' where the derivatives may be rewritten using the up-to relation; 'bisimulation up-to-context' [20], where a common context may be removed from matching derivatives. Further enhancements may exploit the peculiarities of the definition of bisimilarity on certain classes of languages: e.g., the up-to-injective-substitution techniques of the π-calculus [7, 22], techniques for shrinking or enlarging the environment in languages with information hiding mechanisms (e.g., existential types, encryption and decryption constructs [1, 24, 23]), frame equivalence in the psi-calculi [17], or higher-order languages [12, 10]. Lastly, it is important to notice that one often wishes to use *combinations* of up-to techniques. For instance, up-to context alone does not appear to be very useful; its strength comes out in association with other techniques, such as up-to bisimilarity or up-to transitivity.

The main problem with up-to techniques is proving their soundness (i.e. ensuring that any 'bisimulation up-to' is contained in bisimilarity). In particular, the

proofs of complex combinations of techniques can be difficult or, at best, long and tedious. And if one modifies the language or the up-to technique, the entire proof has to be redone from scratch. Indeed the soundness of some up-to techniques is quite fragile, and may break when such variations are made. For instance, in certain models up-to bisimilarity may fail for weak bisimilarity, and in certain languages up-to bisimilarity and context may fail even if bisimilarity is a congruence relation and is strong (treating internal moves as any other move).

This problem has been the motivation for the development of a theory of enhancements, summarised in [18]. Expressed in the general fixed-point theory on complete lattices, this theory has been fully developed for both strong and weak bisimilarity, in the case of first-order labelled transition systems (LTSs) where transitions represent pure synchronisations among processes. In this framework, up-to techniques are represented using *compatible* functions, whose class enjoys nice algebraic properties. This allows one to derive complex up-to techniques algebraically, by composing simpler techniques by means of a few operators.

Only a small part of the theory has been transported onto other forms of transition systems, on a case by case basis. Transferring the whole theory would be a substantial and non-trivial effort. Moreover it might have limited applicability, since this work would probably have to be based on specific shapes for transitions and bisimilarity (a wide range of variations exist, e.g., in higher-order languages).

Here we explore a different approach to the transport of the theory of bisimulation enhancements onto richer languages. The approach consists in exhibiting fully abstract translations of the more sophisticated LTSs and bisimilarities onto first-order LTSs and bisimilarity. This allows us to import directly the existing theory for first-order bisimulation enhancements onto the new languages. Most importantly, the schema allows us to combine up-to techniques for the richer languages. The only additional ingredient that has to be provided manually is the soundness of some up-to techniques that are specific to the new languages. This typically includes the up-to context techniques, since those contexts are not first-order.

Our hope is that the method proposed here will make it possible to obtain a single formalised library about up-to techniques, that can be reused for a wide range of calculi: currently, all existing formalisations of such techniques in a proof assistant are specific to a given calculus: π-calculus [5, 4], the psi-calculi [17], or a miniML language [6].

We consider three languages: the π-calculus, the call-by-name λ-calculus, and an imperative call-by-value λ-calculus. Other calculi like the Higher-Order π-calculus can be handled in a similar way; we omit the details here for lack of space. We moreover focus on weak bisimilarity, since its theory is more delicate than that of strong bisimilarity. When we translate a transition system into a first-order one, the grammar for the labels can be complex (e.g. include terms, or labels, or contexts). What makes the system 'first-order' is that such labels are taken as syntactic atomic objects, that may only be checked for syntactic equality. Note that full abstraction of the translation does not imply that the up-to techniques come for free: further conditions must be ensured. We shall see this with the π-calculus, where early bisimilarity can be handled but not the late one.

Forms of up-to context have already been derived for the languages we consider in this paper [11, 22, 21]. The corresponding soundness proofs are difficult (espe-

cially in λ-calculi), and require a mix of induction (on contexts) and coinduction (to define bisimulations). Recasting up-to context within the theory of bisimulation enhancements has several advantages. First, this allows us to combine this technique with other techniques, directly. Second, substitutivity (or congruence) of bisimilarity becomes a corollary of the compatibility of the up-to-context function (in higher-order languages these two kinds of proofs are usually hard and very similar). And third, the theory allows us to decompose the up-to context function into smaller pieces, essentially one for each operator of the language, yielding more modular proofs, also allowing, if needed, to rule out those contexts that do not preserve bisimilarity (e.g., input prefix in the π-calculus).

The translation of the π-calculus LTS into a first-order LTS follows the schema of abstract machines for the π-calculus (e.g., [25]) in which the issue of the choice of fresh names is resolved by ordering the names. Various forms of bisimulation enhancements have appeared in papers on the π-calculus or dialects of it. A translation of higher-order π-calculi into first-order processes has been proposed by Koutavas et al [8]. While the shape of our translations of λ-calculi is similar, our LTSs differ since they are designed to recover the theory of bisimulation enhancements. In particular, using the LTSs from [8] would lead to technical problems similar to those discussed in Remark 2. In the λ-calculus, limited forms of up-to techniques have been developed for applicative bisimilarity, where the soundness of the up-to context has still open problems [12, 11]. More powerful versions of up-to context exist for forms of bisimilarity on open terms (e.g., open bisimilarity or head-normal-form bisimilarity) [13]. Currently, the form of bisimilarity for closed higher-order terms that allows the richest range of up-to techniques is environmental bisimilarity [21, 9]. However, even in this setting, the proofs of combinations of up-to techniques are usually long and non-trivial. Our translation of higher-order terms to first-order terms is designed to recover environmental bisimilarity.

In Section 6, we show an example of how the wide spectrum of up-to techniques made available via our translations allows us to simplify relations needed in bisimilarity proofs, facilitating their description and reducing their size.

2 First-order bisimulation and up-to techniques

A *first-order Labelled Transition System*, briefly LTS, is a triple $(Pr, Act, \longrightarrow)$ where Pr is a non-empty set of states (or processes), Act is the set of *actions* (or *labels*), and $\longrightarrow \subseteq Pr \times Act \times Pr$ is the *transition relation*. We use P, Q, R to range over the processes of the LTS, and μ to range over the labels in Act, and, as usual, write $P \xrightarrow{\mu} Q$ when $(P, \mu, Q) \in \longrightarrow$. We assume that Act includes a special action τ that represents an internal activity of the processes. We derive bisimulation from the notion of *progression* between relations.

Definition 1. *Suppose $\mathcal{R}, \mathcal{S}$ are relations on the processes of an LTS. Then $\mathcal{R}$* strongly progresses to *$\mathcal{S}$, written $\mathcal{R} \rightsquigarrow_{\mathbf{sp}} \mathcal{S}$, if $\mathcal{R} \subseteq \mathcal{S}$ and if $P \mathcal{R} Q$ implies:*

- *whenever $P \xrightarrow{\mu} P'$ there is Q' s.t. $Q \xrightarrow{\mu} Q'$ and $P' \mathcal{S} Q'$;*
- *whenever $Q \xrightarrow{\mu} Q'$ there is P' s.t. $P \xrightarrow{\mu} P'$ and $P' \mathcal{S} Q'$.*

A relation $\mathcal{R}$ is a strong bisimulation *if $\mathcal{R} \rightsquigarrow_{\mathbf{sp}} \mathcal{R}$; and* strong bisimilarity, $\sim$, *is the union of all strong bisimulations.*

To define weak progression we need weak transitions, defined as usual: $P \xrightarrow{\hat{\mu}} P'$ means $P \xrightarrow{\mu} P'$ or $\mu = \tau$ and $P = P'$; and $\stackrel{\hat{\mu}}{\Longrightarrow}$ is $\Longrightarrow\xrightarrow{\hat{\mu}}\Longrightarrow$ where $\Longrightarrow$ is the reflexive transitive closure of $\xrightarrow{\tau}$. *Weak progression*, $\mathcal{R} \rightsquigarrow_{\mathbf{wp}} \mathcal{S}$, and *weak bisimilarity*, $\approx$, are obtained from Definition 1 by allowing the processes to answer using $\stackrel{\hat{\mu}}{\Longrightarrow}$ rather than $\xrightarrow{\mu}$.

Below we summarise the ingredients of the theory of bisimulation enhancements for first-order LTSs from [18] that will be needed in the sequel. We use f and g to range over functions on relations. Each such function represents a potential up-to technique; only the *sound* functions, however, qualify as up-to techniques:

Definition 2. *A function f is* sound for $\sim$ *if $\mathcal{R} \rightsquigarrow_{\mathbf{sp}} f(\mathcal{R})$ implies $\mathcal{R} \subseteq \sim$, for all $\mathcal{R}$; similarly, f is* sound for $\approx$ *if $\mathcal{R} \rightsquigarrow_{\mathbf{wp}} f(\mathcal{R})$ implies $\mathcal{R} \subseteq \approx$, for all $\mathcal{R}$.*

Unfortunately, the class of sound functions does not enjoy good algebraic properties. As a remedy to this, the subset of *compatible* functions has been proposed. The concepts in the remainder of the section can be instantiated with both strong and weak bisimilarities; we thus use $\mathbf{p}$ to range over $\mathbf{sp}$ or $\mathbf{wp}$.

Definition 3. *We write $f \rightsquigarrow_{\mathbf{p}} g$ when $\mathcal{R} \rightsquigarrow_{\mathbf{p}} \mathcal{S}$ implies $f(\mathcal{R}) \rightsquigarrow_{\mathbf{p}} g(\mathcal{S})$ for all $\mathcal{R}$ and $\mathcal{S}$. A monotone function f on relations is* $\mathbf{p}$-compatible *if $f \rightsquigarrow_{\mathbf{p}} f$.*

In other terms [18], f is $\mathbf{p}$-compatible iff $f \circ \mathbf{p} \subseteq \mathbf{p} \circ f$ where $\mathbf{p}(\mathcal{S})$ is the union of all $\mathcal{R}$ such that $\mathcal{R} \rightsquigarrow_{\mathbf{p}} \mathcal{S}$ and $\circ$ denotes function composition. Note that $\mathcal{R} \rightsquigarrow_{\mathbf{p}} \mathcal{S}$ is equivalent to $\mathcal{R} \subseteq \mathbf{p}(\mathcal{S})$.

Lemma 1. *If f is $\mathbf{sp}$-compatible, then f is sound for $\sim$; if f is $\mathbf{wp}$-compatible, then f is sound for $\approx$.*

Simple examples of compatible functions are the identity function and the function mapping any relation onto bisimilarity (for the strong or weak case, respectively). The class of compatible functions is closed under function composition and union (where the union $\cup F$ of a set of functions F is the point-wise union mapping $\mathcal{R}$ to $\bigcup_{f \in F} f(\mathcal{R})$), and thus under omega-iteration (where the omega-iteration f^ω of a function f maps $\mathcal{R}$ to $\bigcup_{n \in \mathbb{N}} f^n(\mathcal{R})$).

Other examples of compatible functions are typically contextual closure functions, mapping a relation into its closure w.r.t. a given set of contexts. For such functions, the following lemma shows that the compatibility of up-to-context implies substitutivity of (strong or weak) bisimilarity.

Lemma 2. *If f is $\mathbf{sp}$-compatible, then $f(\sim) \subseteq \sim$; similarly if f is $\mathbf{wp}$-compatible, then $f(\approx) \subseteq \approx$.*

Certain closure properties for compatible functions however only hold in the strong case. The main example is the *chaining operator* $\frown$, which implements relational composition:

$$f \frown g\,(\mathcal{R}) \triangleq f(\mathcal{R})\; g(\mathcal{R})$$

where $f(\mathcal{R})$ $g(\mathcal{R})$ indicates the composition of the two relations $f(\mathcal{R})$ and $g(\mathcal{R})$. Using chaining we can obtain the compatibility of the function 'up to transitivity' mapping any relation $\mathcal{R}$ onto its reflexive and transitive closure $\mathcal{R}^\star$. Another example of **sp**-compatible function is 'up to bisimilarity' ($\mathcal{R} \mapsto \sim\mathcal{R}\sim$).

In contrast, in the weak case bisimulation up to bisimilarity is unsound. This is a major drawback in up-to techniques for weak bisimilarity, which can be partially overcome by resorting to the *expansion* relation $\gtrsim$ [3]. Expansion is an asymmetric refinement of weak bisimilarity whereby $P \gtrsim Q$ holds if P and Q are bisimilar and, in addition, Q is at least as efficient as P, in the sense that Q is capable of producing the same activity as P without ever performing more internal activities (the τ-actions); see Appendix A for its definition. Up-to-expansion yields a function ($\mathcal{R} \mapsto \gtrsim\mathcal{R}\lesssim$) that is **wp**-compatible. As a consequence, the same holds for the 'up-to expansion and contexts' function. More sophisticated up-to techniques can be obtained by carefully adjusting the interplay between visible and internal transitions, and by taking into account termination hypotheses [18].

Some further compatible functions are the functions **sp** and **wp** themselves (indeed a function f is **p**-compatible if $f \circ \mathbf{p} \subseteq \mathbf{p} \circ f$, hence trivially f can be replaced by **p** itself). Intuitively, the use of **sp** and **wp** as up-to techniques means that, in a diagram-chasing argument, the two derivatives need not be related; it is sufficient that the derivatives of such derivatives be related. Accordingly, we sometimes call functions **sp** and **wp** *unfolding* functions. We will use **sp** in the example in Section 6 and **wp** in Sections 4 and 5, when proving the **wp**-compatibility of the up to context techniques.

Last, note that to use a function f in combinations of up-to techniques, it is actually not necessary that f be **p**-compatible: for example proving that f progresses to $f \cup g$ and g progresses to g is enough, as then $f \cup g$ would be compatible. Extending this reasoning allows us to make use of 'second-order up-to techniques' to reason about compatibility of functions. When F is a set of functions, we say that F is **p**-*compatible up to* if for all f in F, it holds that $f \rightsquigarrow_{\mathbf{p}} (g \cup (\cup F))^\omega$ for a function g that has already been proven compatible. (We sometimes say that F is **p**-*compatible up to* g, to specify which compatible function is employed.) Lemma 1 and 2 remain valid when 'compatible' is replaced by 'compatible up to'.

Terminology We will simply say that a function is *compatible* to mean that it is both **sp**-compatible and **wp**-compatible; similarly for compatibility up to. In languages defined from a grammar, a context C is a term with numbered holes $[\cdot]_1, \ldots, [\cdot]_n$, and each hole $[\cdot]_i$ can appear any number of times in C.

3 The π-calculus

The syntax and operational semantics of the π-calculus are recalled in Appendix B. We consider the early transition system, in which transitions are of the forms

$$P \overset{ab}{\longmapsto}_\pi P' \qquad P \overset{\overline{a}b}{\longmapsto}_\pi P' \qquad P \overset{\overline{a}(b)}{\longmapsto}_\pi P' \ .$$

In the third transition, called bound output transition, name b is a binder for the free occurrences of b in P' and, as such, it is subject to α-conversion. The definition of bisimilarity takes α-conversion into account. The clause for bound output of strong early bisimilarity says ($\mathrm{fn}(Q)$ indicates the names free in Q):

– if $P \overset{\overline{a}(b)}{\longmapsto}_\pi P'$ and $b \notin \mathrm{fn}(Q)$ then $Q \overset{\overline{a}(b)}{\longmapsto}_\pi Q'$ for some Q' such that $P' \sim Q'$.

(The complete definition of bisimilarity is recalled in Appendix B). When translating the π-calculus semantics to a first-order one, α-conversion and the condition $b \notin \mathrm{fn}(Q)$ have to be removed. To this end, one has to force an agreement between two bisimilar process on the choice of the bound names appearing in transitions. We obtain this by considering *named processes* (c, P) in which c is bigger or equal to all names in P. For this to make sense we assume an enumeration of the names and use $\leq$ as the underlying order, and $c+1$ for name following c in the enumeration; for a set of names N, we also write $c \geq N$ to mean $c \geq a$ for all $a \in N$.

The rules below define the translation of the π-calculus transition system to a first-order LTS. In the first-order LTS, the grammar for labels is the same as that of the original LTS; however, for a named process (c, P) the only name that may be exported in a bound output is $c+1$; similarly only names that are below or equal to $c+1$ may be imported in an input transition. (Indeed, testing for all fresh names $b > c$ is unnecessary, doing it only for one ($b = c+1$) is enough.) This makes it possible to use the ordinary definition of bisimilarity for first-order LTS, and thus recover the early bisimilarity on the source terms.

$$\frac{P \overset{\tau}{\longmapsto}_\pi P'}{(c,P) \overset{\tau}{\longrightarrow} (c,P')} \qquad \frac{P \overset{ab}{\longmapsto}_\pi P'}{(c,P) \overset{ab}{\longrightarrow} (c,P')}\ b \leq c \qquad \frac{P \overset{\overline{a}b}{\longmapsto}_\pi P'}{(c,P) \overset{\overline{a}b}{\longrightarrow} (c,P')}\ b \leq c$$

$$\frac{P \overset{ab}{\longmapsto}_\pi P'}{(c,P) \overset{ab}{\longrightarrow} (b,P')}\ b = c+1 \qquad \frac{P \overset{\overline{a}(b)}{\longmapsto}_\pi P'}{(c,P) \overset{\overline{a}(b)}{\longrightarrow} (b,P')}\ b = c+1$$

We write π^1 for the first-order LTS derived from the above translation of the π-calculus. Although the labels of the source and target transitions have a similar shape, the LTS in π^1 is first-order because labels are taken as purely syntactic objects (without α-conversion). We write $\sim^{\mathrm{e}}$ and $\approx^{\mathrm{e}}$ for strong and weak early bisimilarity of the π-calculus.

Theorem 1. *Assume $c \geq \mathrm{fn}(P) \cup \mathrm{fn}(Q)$. Then we have: $P \sim^{\mathrm{e}} Q$ iff $(c,P) \sim (c,Q)$, and $P \approx^{\mathrm{e}} Q$ iff $(c,P) \approx (c,Q)$.*

The above full abstraction result allows us to import the theory of up-to techniques for first-order LTSs and bisimilarity, both in the strong and the weak case. We have however to prove the soundness of up-to techniques that are specific to the π-calculus. Function isub implements 'up to injective name substitutions':

$$\mathsf{isub}(\mathcal{R}) \triangleq \{((d,P\sigma),(d,Q\sigma)) \ \text{ s.t. } (c,P)\ \mathcal{R}\ (c,Q),\ \mathrm{fn}(P\sigma) \cup \mathrm{fn}(Q\sigma) \leq d, \text{ and } \sigma \text{ is injective on } \mathrm{fn}(P) \cup \mathrm{fn}(Q) \ \}\ .$$

Another function for manipulating names, str, allows us to replace the index c in a named process (c, P) with a lower one:

$$\mathsf{str}(\mathcal{R}) \triangleq \{((d,P),(d,Q)) \ \text{ s.t. } (c,P)\ \mathcal{R}\ (c,Q) \text{ and } \mathrm{fn}(P,Q) \leq d\ \}\ .$$

Lemma 3. *The set* $\{\mathsf{isub}, \mathsf{str}\}$ *is compatible up to.*

The up-to-context function is decomposed into a set of smaller context functions, called *initial* [18], one for each operator of the π-calculus. The only exception to this is the input prefix, since early bisimilarity in the π-calculus is not preserved by this operator. We write $\mathcal{C}_o$, $\mathcal{C}_\nu$, $\mathcal{C}_!$, $\mathcal{C}_|$, and $\mathcal{C}_+$ for these initial context functions, respectively returning the closure of a relation under the operators of output prefix, restriction, replication, parallel composition, and sum.

Definition 4. *If* $\mathcal{R}$ *is a relation on* π^1*, we define* $\mathcal{C}_o(\mathcal{R})$*,* $\mathcal{C}_\nu(\mathcal{R})$*,* $\mathcal{C}_!(\mathcal{R})$*,* $\mathcal{C}_|(\mathcal{R})$ *and* $\mathcal{C}_+(\mathcal{R})$ *by saying that whenever* $(c, P)\ \mathcal{R}\ (c, Q)$*,*

- $(c, \overline{a}b.P)\ \ \mathcal{C}_o(\mathcal{R})\ \ (c, \overline{a}b.Q)$*, for any* a, b *with* $a, b \leq n$*,*
- $(c, \nu a.P)\ \ \mathcal{C}_\nu(\mathcal{R})\ \ (c, \nu a.Q)$*,*
- $(c, !P)\ \ \mathcal{C}_!(\mathcal{R})\ \ (c, !Q)$*;*

and, whenever $(c, P_1)\ \mathcal{R}\ (c, Q_1)$ *and* $(c, P_2)\ \mathcal{R}\ (c, Q_2)$*,*

- $(c, P_1 \mid Q_1)\ \ \mathcal{C}_|(\mathcal{R})\ \ (c, P_2 \mid Q_2)$*,*
- $(c, P_1 + Q_1)\ \ \mathcal{C}_+(\mathcal{R})\ \ (c, P_2 + Q_2)$*.*

While bisimilarity in the π-calculus is not preserved by input prefix, a weaker rule holds (where $=$ can be $\sim^{\mathrm{e}}$ or $\approx^{\mathrm{e}}$):

$$\frac{P = Q \quad \text{and } P\{{}^{c}\!/\!_{b}\} = Q\{{}^{c}\!/\!_{b}\} \text{ for each } c \text{ free in } P, Q}{a(b).P = a(b).Q} \tag{1}$$

We define $\mathcal{C}_i$, the function for input prefix, accordingly: we have $(d, a(b).P)\ \mathcal{C}_i(\mathcal{R})$ $(d, a(b).Q)$ if $a \leq d$ and $(d+1, P\{c/b\})\ \mathcal{R}\ (d+1, Q\{c/b\})$ for all $c \leq d+1$.

Theorem 2. *The set* $\{\mathcal{C}_o, \mathcal{C}_i, \mathcal{C}_\nu, \mathcal{C}_!, \mathcal{C}_|, \mathcal{C}_+\}$ *is* $\mathbf{sp}$*-compatible up to* $\mathsf{isub} \cup \mathsf{str}$*.*

Weak bisimilarity is not preserved by sums, only by guarded sums, whose function is $\mathcal{C}_{g+} \triangleq \mathcal{C}_+^\omega \circ (\mathcal{C}_o \cup \mathcal{C}_i)$.

Theorem 3. *The set* $\{\mathcal{C}_o, \mathcal{C}_i, \mathcal{C}_\nu, \mathcal{C}_!, \mathcal{C}_|, \mathcal{C}_{g+}\}$ *is* $\mathbf{wp}$*-compatible up to* $\mathsf{isub} \cup \mathsf{str}$*.*

As a byproduct of the compatibility of these initial context functions, and using Lemma 2, we derive the standard substitutivity properties of strong and weak early bisimilarity, including the rule (1) for input prefix.

Corollary 1. *In the* π*-calculus, relations* $\sim^{\mathrm{e}}$ *and* $\approx^{\mathrm{e}}$ *are preserved by the operators of output prefix, replication, parallel composition, restriction;* $\sim^{\mathrm{e}}$ *is also preserved by sum, whereas* $\approx^{\mathrm{e}}$ *is only preserved by guarded sums. Moreover, rule (1) is valid both for* $\sim^{\mathrm{e}}$ *and* $\approx^{\mathrm{e}}$*.*

Remark 1. Late bisimilarity makes use of transitions $P \overset{a(b)}{\longmapsto}_\pi P'$ where b is bound, the definition of bisimulation containing a quantification over names. To capture this bisimilarity in a first-order LTS we would need to have two transitions for the input $a(b)$: one to fire the input a, leaving b uninstantiated, and another to instantiate b. While such a translation does yield full abstraction for both strong and weak late bisimilarities, the decomposition of an input transition into two steps prevents us from obtaining the compatibility of up to context.

4 Call-by-name λ-calculus

To study the applicability of our approach to higher-order languages, we investigate the pure call-by-name λ-calculus, referred to as ΛN in the sequel.

We use M, N to range over the set Λ of λ-terms, and x, y, z to range over variables. The standard syntax of λ-terms, and the rules for call-by-name reduction, are recalled in Appendix C. We assume the familiar concepts of free and bound variables and substitutions, and identify α-convertible terms. The only values are the λ-abstractions $\lambda x.M$. In this section and in the following one, results and definitions are presented on closed terms; extension to open terms is made using closing abstractions (i.e., abstracting on all free variables). The reduction relation of ΛN is $\longmapsto_{\mathsf{n}}$, and $\Longmapsto_{\mathsf{n}}$ its reflexive and transitive closure.

As bisimilarity for the λ-calculus we consider *environmental bisimilarity* [21, 9], which allows a set of up-to techniques richer than Abramsky's applicative bisimilarity [2], even if the two notions actually coincide, together with contextual equivalence. Environmental bisimilarity makes a clear distinction between the tested terms and the environment. An element of an environmental bisimulation has, in addition to the tested terms M and N, a further component $\mathcal{E}$, the environment, which expresses the observer's current knowledge. When an input from the observer is required, the arguments supplied are terms that the observer can build using the current knowledge; that is, terms obtained by composing the values in $\mathcal{E}$ using the operators of the calculus. An *environmental relation* is a set of elements each of which is of the form $(\mathcal{E}, M, N)$ or $\mathcal{E}$, and where M, N are closed terms and $\mathcal{E}$ is a relation on closed values. We use $\mathcal{X}, \mathcal{Y}$ to range over environmental relations. In a triple $(\mathcal{E}, M, N)$ the relation component $\mathcal{E}$ is the *environment*, and M, N are the *tested terms*. We write $M\ \mathcal{X}_{\mathcal{E}}\ N$ for $(\mathcal{E}, M, N) \in \mathcal{X}$. We write $\mathcal{E}^\star$ for the closure of $\mathcal{E}$ under contexts. We only define the weak version of the bisimilarity; its strong version is obtained in the expected way.

Definition 5. *An environmental relation $\mathcal{X}$ is an* environmental bisimulation *if*

1. *$M\ \mathcal{X}_{\mathcal{E}}\ N$ implies:*
   (a) if $M \longmapsto_{\mathsf{n}} M'$ then $N \Longmapsto_{\mathsf{n}} N'$ and $M'\ \mathcal{X}_{\mathcal{E}}\ N'$;
   (b) if $M = V$ then $N \Longmapsto_{\mathsf{n}} W$ and $\mathcal{E} \cup \{(V, W)\} \in \mathcal{X}$;
   (c) the converse of the above two conditions, on N;
2. *if $\mathcal{E} \in \mathcal{X}$ then for all $(\lambda x.P, \lambda x.Q) \in \mathcal{E}$ and for all $(M, N) \in \mathcal{E}^\star$ it holds that $P\{^M/_x\}\ \mathcal{X}_{\mathcal{E}}\ Q\{^N/_x\}$.*

Environmental bisimilarity, *$\approx^{env}$, is the union of all environmental bisimulations.*

For the translation of environmental bisimilarity to first-order, a few issues have to be resolved. For instance, an environmental bisimilarity contains both triples $(\mathcal{E}, M, N)$, and pure environments $\mathcal{E}$, which shows up in the difference between clauses (1) and (2) of Definition 5. Moreover, the input supplied to tested terms may be constructed using arbitrary contexts.

We write ΛN^1 for the first-order LTS resulting from the translation of ΛN. The states of ΛN^1 are sequences of λ-terms in which only the last one need not be a value. We use Γ and Δ to range over sequences of values only; thus (Γ, M) indicates a sequence of λ-values followed by M; and Γ_i is the i-th element in Γ.

For an environment $\mathcal{E}$, we write $\mathcal{E}_1$ for an ordered projection of the pairs in $\mathcal{E}$ on the first component, and $\mathcal{E}_2$ is the corresponding projection on the second component. In the translation, intuitively, a triple $(\mathcal{E}, M, N)$ of an environmental bisimulation is split into the two components $(\mathcal{E}_1, M)$ and $(\mathcal{E}_2, N)$. Similarly, an environment $\mathcal{E}$ is split into $\mathcal{E}_1$ and $\mathcal{E}_2$. We write $C[\Gamma]$ for the term obtained by replacing each hole $[\cdot]_1$ in C with the value Γ_i. The rules for transitions in ΛN^1 are as follows:

$$\frac{M \longmapsto_{\mathsf{n}} M'}{(\Gamma, M) \xrightarrow{\tau} (\Gamma, M')} \qquad \frac{\Gamma_i(C[\Gamma]) \longmapsto_{\mathsf{n}} M'}{\Gamma \xrightarrow{i,C} (\Gamma, M')} \tag{2}$$

The first rule says that if M reduces to M' in ΛN then M can also reduce in ΛN^1, in any environment. The second rule implements the observations in clause (2) of Definition 5: in an environment Γ (only containing values), any component Γ_i can be tested by supplying, as input, a term obtained by filling a context C with values from Γ itself. The label of the transition records the position i and the context chosen. As the rules show, the labels of ΛN^1 include the special label τ, and can also be of the form i, C where i is a integer and C a context.

Theorem 4. *$M \approx^{\mathsf{env}}_{\mathcal{E}} N$ iff $(\mathcal{E}_1, M) \approx (\mathcal{E}_2, N)$.*

(The theorem also holds for the strong versions of the bisimilarities.) Again, having established full abstraction with respect to a first-order transition system and ordinary bisimilarity, we can inherit the theory of bisimulation enhancements. We have however to check up-to techniques that are specific to environmental bisimilarity. A useful such technique is 'up to environment', which allows us to replace an environment with a larger one; $\mathsf{w}(\mathcal{R})$ is the smallest relation that includes $\mathcal{R}$ and such that, whenever $(V, \Gamma, M)\ \mathsf{w}(\mathcal{R})\ (W, \Delta, N)$ then also $(\Gamma, M)\ \mathsf{w}(\mathcal{R})\ (\Delta, N)$, where V and W are any values. (Here w stands for 'weakening' as, from Lemmas 2 and 4, if $(V, \Gamma, M) \approx (W, \Delta, N)$ then $(\Gamma, M) \approx (\Delta, N)$.)

Lemma 4. *Function w is compatible.*

Somehow dual to weakening is the strengthening of the environment, in which a component of an environment can be removed. However this is only possible if the component removed is 'redundant', that is, it can be obtained by gluing other pieces of the environment within a context; strengthening is captured by the following str function: $(\Gamma, C_v[\Gamma], M)\ \mathsf{str}(\mathcal{R})\ (\Delta, C_v[\Delta], N)$ whenever $(\Gamma, M)\ \mathcal{R}\ (\Delta, N)$ and C_v is a value context (i.e., the outermost operator is an abstraction). We derive the compatibility up to of str in Theorem 5.

For up-to context, we need to distinguish between arbitrary contexts and evaluation contexts. There are indeed substitutivity properties, and corresponding up-to techniques, that only hold for the latter contexts. A hole $[\cdot]_i$ of a context C is in a *redex position* if the context obtained by filling all the holes but $[\cdot]_i$ with values is an evaluation context. Below C ranges over arbitrary contexts, whereas E ranges over contexts whose first hole is in redex position.

$$\begin{aligned} \mathcal{C}(\mathcal{R}) &\triangleq \big\{((\Gamma, C[\Gamma]), (\Delta, C[\Delta])) \qquad \text{s.t. } \Gamma\ \mathcal{R}\ \Delta\big\} \\ \mathcal{C}_{\mathsf{e}}(\mathcal{R}) &\triangleq \big\{((\Gamma, E[M, \Gamma]), (\Delta, E[N, \Delta])) \text{ s.t. } (\Gamma, M)\ \mathcal{R}\ (\Delta, N)\big\} \end{aligned}$$

Theorem 5. *The set* $\{\mathsf{str}, \mathcal{C}, \mathcal{C}_e\}$ *is compatible up to.*[3]

For the proof, we establish the progression property separately for each function in $\{\mathsf{str}, \mathcal{C}, \mathcal{C}_\mathsf{e}\}$, using simple diagram-chasing arguments (together with induction on the structure of a context). Once more, the compatibility of the up to context functions entails also the substitutivity properties of environmental bisimilarity. In [21] the two aspects (substitutivity and up-to context) had to be proved separately, with similar proofs. Moreover the two cases of contexts (arbitrary contexts and evaluation contexts) had to be considered at the same time, within the same proof. Here, in contrast, the machinery of compatible function allows us to split the proof into two simpler proofs.

Remark 2. A transition system ensuring full abstraction as in Theorem 4 does not guarantee the compatibility of the up-to techniques specific to the language in consideration. For instance, a simpler and maybe more natural alternative to the second transition in (2) is the following one:

$$\frac{}{\Gamma \xrightarrow{i,C} (\Gamma, \Gamma_i(C[\Gamma]))} \tag{3}$$

With this rule, full abstraction holds, but up-to context is unsound: for any Γ and Δ, the singleton relation $\{(\Gamma, \Delta)\}$ is a bisimulation up to $\mathcal{C}$: indeed, using rule (3), the derivatives of the pair Γ, Δ are of the shape $\Gamma_i(C[\Gamma])$, $\Delta_i(C[\Delta])$, and they can be discarded immediately, up to the context $[\cdot]_i C$. If up-to context were sound then we would deduce that any two terms are bisimilar. (The rule in (2) prevents such a behaviour since it ensures that the tested values are 'consumed' immediately.)

5 Imperative call-by-value λ-calculus

In this section we study the addition of imperative features (higher-order references, that we call locations), to a call-by-value λ-calculus. It is known that finding powerful reasoning techniques for imperative higher-order languages is a hard problem. The language, ΛR, is a simplified variant of that in [10, 21]. The syntax of terms, values, and evaluation contexts, as well as the reduction semantics are given in Figure 1. A λ-term M is run in a *store*: a partial function from locations to closed values, whose domain includes all free locations of both M and its own co-domain. We use letters s, t to range over stores. New store locations may be created using the operator $\nu\ell\, M$; the content of a store location ℓ may be rewritten using $\mathsf{set}_\ell V$, or read using $\mathsf{get}_\ell V$ (the former instruction returns a value, namely the identity $I \triangleq \lambda x.x$, and the argument of the latter one is ignored). We denote the reflexive and transitive closure of $\longmapsto_\mathsf{R}$ by $\Longmapsto_\mathsf{R}$.

Note that in contrast with the languages in [10, 21], locations are not directly first-class values values; the expressive power is however the same: a first-class location ℓ can always be encoded as the pair $(\mathsf{get}_\ell, \mathsf{set}_\ell)$.

We present the first-order LTS for ΛR, and then we relate the resulting strong and weak bisimilarities directly with contextual equivalence (the reference equivalence in λ-calculi). Alternatively, we could have related the first-order bisimilarities

[3] in the weak case, **wp**-compatible up to $\mathbf{wp} \cup \mathbf{e}$ where $\mathbf{e}$ is 'up to expansion'.

$$M ::= x \mid MM \mid \nu\ell\, M \mid V \qquad V ::= \lambda x.M \mid \mathsf{set}_\ell \mid \mathsf{get}_\ell \qquad E ::= [\cdot] \mid EV \mid ME$$

$$\frac{}{(s;(\lambda x.M)V) \longmapsto_{\mathsf{R}} (s;M\{V/x\})} \qquad \frac{\ell \notin \mathsf{dom}(s)}{(s;\nu\ell\, M) \longmapsto_{\mathsf{R}} (s[\ell \mapsto I];M)}$$

$$\frac{\ell \in \mathsf{dom}(s)}{(s;\mathsf{get}_\ell V) \longmapsto_{\mathsf{R}} (s;s[\ell])} \qquad \frac{\ell \in \mathsf{dom}(s)}{(s;\mathsf{set}_\ell V) \longmapsto_{\mathsf{R}} (s[\ell \mapsto V];I)} \qquad \frac{(s;M) \longmapsto_{\mathsf{R}} (s';M')}{(s;E[M]) \longmapsto_{\mathsf{R}} (s';E[M'])}$$

Fig. 1. The imperative λ-calculus

to the environmental bisimilarities of ΛR, and then inferred the correspondence with contextual equivalence from known results about environmental bisimilarity, as we did for ΛN.

We write $(s;M)\downarrow$ when M is a value; and $(s;M)\Downarrow$ if $(s;M) \Longrightarrow_{\mathsf{R}}\downarrow$. For the definition of contextual equivalence, we distinguish the cases of values and of arbitrary terms, because they have different substitutivity properties: values can be tested in arbitrary contexts, while arbitrary terms must be tested only in evaluation contexts. As in [21], we consider contexts that do not contain free locations (they can contain bound locations). We refer to [21] for more details on these aspects.

Definition 6. – *For values V, W, we write $(s;V) \equiv (t;W)$ when $(s;C[V])\Downarrow$ iff $(t;C[W])\Downarrow$, for all location-free context C.*
– *For terms M and N, we write $(s;M) \equiv (t;N)$ when $(s;E[M])\Downarrow$ iff $(t;E[N])\Downarrow$, for all location-free evaluation context E.*

We now define ΛR^1, the first-order LTS for ΛR. The states and the transitions for ΛR^1 are similar to those for the pure λ-calculus of Section 4, with the addition of a component for the store. The two transitions (2) of call-by-name λ-calculus become:

$$\frac{(s;M) \longmapsto_{\mathsf{R}} (s';M')}{(s;\Gamma,M) \xrightarrow{\tau} (s';\Gamma,M')} \qquad \frac{\Gamma' = \Gamma,\mathsf{getset}(r) \quad (s \uplus r[\Gamma'];\Gamma_i(C[\Gamma'])) \longmapsto_{\mathsf{R}} (s';M')}{(s;\Gamma) \xrightarrow{i,C,\mathtt{cod}(r)} (s';\Gamma',M')}$$

The first rule is the analogous of the first rule in (2). The important differences are on the second rule. First, since we are *call-by-value*, C now ranges over $\mathbb{C}_v$, the set of *value contexts* (i.e., contexts of the form $\lambda x.C'$) without free locations. Moreover, since we are now *imperative*, in a transition we must permit the creation of new locations, and a term supplied by the environment should be allowed to use them. In the rule, the new store is represented by r (whose domain has to be disjoint from that of s) Correspondingly, to allow manipulation of these locations from the observer, for each new location ℓ we make set_ℓ and get_ℓ available, as an extension of the environment; in the rule, these are collectively written $\mathsf{getset}(r)$, and Γ' is the extended environment. Finally, we must initialise the new store, using terms that are created out of the extended environment Γ'; that is, each new location ℓ is initialised with a term $D_\ell[\Gamma']$ (for $D_\ell \in \mathbb{C}_v$). Moreover, the contexts D_ℓ chosen

must be made visible in the label of the transition. To take care of these aspects, we view r as a *store context*, a tuple of assignments $\ell \mapsto D_\ell$. Thus the initialisation of the new locations is written $r[\Gamma']$; and, denoting by $\mathtt{cod}(r)$ the tuple of the contexts D_ℓ in r, we add $\mathtt{cod}(r)$ to the label of the transition. Note also that, although C and D_ℓ are location-free, their holes may be instantiated with terms involving the set_ℓ and get_ℓ operators, and these allow manipulation of the store.

Once more, on the (strong and weak) bisimilarities that are derived from this first-order LTS we can import the theory of compatible functions and bisimulation enhancements. Concerning additional up-to functions, specific to ΛR, the functions w, str, $\mathcal{C}$ and $\mathcal{C}_\mathsf{e}$ are adapted from Section 4 in the expected manner—contexts C_v, C and E must be location-free. A further function for ΛR is store, which manipulates the store by removing locations that do not appear elsewhere (akin to garbage collection); thus, $\mathsf{store}(\mathcal{R})$ is the set of all pairs

$$((s \uplus r[\Gamma']; \Gamma', M),\ (t \uplus r[\Delta']; \Delta', N))$$

such that $(s; \Gamma, N)\ \mathcal{R}\ (t; \Delta, N)$, and with $\Gamma' = \Gamma, \mathsf{getset}(r)$ and $\Delta' = \Delta, \mathsf{getset}(r)$.

Lemma 5. *The set* $\{\mathsf{w}, \mathsf{str}, \mathcal{C}_\mathsf{e}, \mathsf{store}, \mathcal{C}\}$ *is compatible up to.*[4]

The techniques $\mathcal{C}$ and $\mathcal{C}_\mathsf{e}$ allow substitutivity under location-free contexts, from which we can derive the soundness part of Theorem 6.

Theorem 6. $(s; M) \approx (t; N)$ *iff* $(s; M) \equiv (t; N)$.

Proof (sketch). Soundness ($\Rightarrow$) follows from congruence by $\mathcal{C}_\mathsf{e}$ (Lemmas 5 and 2) and completeness ($\Leftarrow$) is obtained by standard means. See Appendix D for details.

Note that substitutivity of bisimilarity is restricted either to values ($\mathcal{C}$), or to evaluation contexts ($\mathcal{C}_\mathsf{e}$). The following lemma provides a sufficient condition for a given law between arbitrary terms to be preserved by arbitrary contexts.

Lemma 6. *Let* $\asymp$ *be any of the relations* $\sim, \approx$, *and* $\gtrsim$. *Suppose* L, R *are* ΛR *terms with* $(s; \Gamma, L) \asymp (s; \Gamma, R)$ *for all environments* Γ *and stores* s. *Then also* $(s; \Gamma, C[L]) \asymp (s; \Gamma, C[R])$, *for any store* s, *environment* Γ *and context* C.

Proof (sketch). We first prove a simplified result in which C is an evaluation context, using techniques $\mathcal{C}_\mathsf{e}$ and store. We then exploit this partial result together with up-to expansion to derive the general result. See Appendix D for more details.

We use this lemma at various places in the example we cover in Section 6. For instance we use it to replace a term $N_1 \triangleq (\lambda x.E[x])M$ (with E an evaluation context) with $N_2 \triangleq E[M]$, under an arbitrary context. Such a property is delicate to prove, even for closed terms, because the evaluation of M could involve reading from a location of the store that itself could contain occurrences of N_1 and N_2.

[4] For strong. For weak, it is **wp**-compatible up to $\mathbf{wp} \cup \mathbf{e}$.

6 An example

We conclude by discussing an example from [10]. It consists in proving a law between terms of ΛR extended with integers, operators for integer addition and subtraction, and a conditional—those constructs are straightforward to accommodate in the presented framework. For readability, we also use the standard notation for store assignment, dereferencing and sequence: $(\ell := M) \triangleq \mathsf{set}_\ell M$, $!\ell \triangleq \mathsf{get}_\ell I$, and $M; N \triangleq (\lambda x.N)M$ where x does not appear in N. The two terms are the following ones:

- $M \triangleq \lambda g.\nu \ell\ \ell := 0; g(\mathsf{incr}_\ell); \mathtt{if}\ !\ell\ \mathtt{mod}\ 2 = 0\ \mathtt{then}\ I\ \mathtt{else}\ \Omega$
- $N \triangleq \lambda g.g(F); I$,

where $\mathsf{incr}_\ell \triangleq \lambda z.\ell := !\ell + 2$, and $F \triangleq \lambda z.I$. Intuitively, those two terms are weakly bisimilar because the location bound by ℓ in the first term will always contain an even number.

This example is also considered in [21] where it is however modified to fit the up-to techniques considered in that paper. The latter are less powerful than those available here thanks to the theory of up-to techniques for first-order LTSs (e.g., up to expansion is not considered in [21]—its addition to environmental bisimulations is non-trivial, having stores and environments as parameters).

We consider two proofs of the example. In comparison with the proof in [21]: (i) we handle the original example from [10], and (ii) the availability of a broader set of up-to techniques and the possibility of freely combining them allows us to work with smaller relations. In the first proof we work up to the store (through the function store) and up to expansion—two techniques that are not available in [21]. In the second proof we exploit the up-to-transitivity technique of Section 2, which is only sound for strong bisimilarity, to further reduce the size of the relation we work with.

First proof. We first employ Lemma 6 to reach a variant similar to that of [21]: we make a 'thunk' out of the test in M, and we make N look similar. More precisely, let $\mathsf{test}_\ell \triangleq \lambda z.\mathtt{if}\ !\ell\ \mathtt{mod}\ 2 = 0\ \mathtt{then}\ I\ \mathtt{else}\ \Omega$, we first prove that

- $M \approx M' \triangleq \lambda g.\nu \ell\ \ell := 0; g(\mathsf{incr}_\ell); \mathsf{test}_\ell I$, and
- $N \approx N' \triangleq \lambda g.g(F); FI$.

It then suffices to prove that $M' \approx N'$, which we do using the following relation:

$$\mathcal{R} \triangleq \left\{ \left(s, M', (\mathsf{incr}_\ell, \mathsf{test}_\ell)_{\ell \in \tilde{\ell}}\right), \left(\emptyset, N', (F, F)_{\ell \in \tilde{\ell}}\right) \ \text{ s.t. } \forall \ell \in \tilde{\ell}, \quad s(\ell) \text{ is even} \right\} \ .$$

The initial pair of terms is generalised by adding any number of private locations, since M' can use itself to create more of them. Relation $\mathcal{R}$ is a weak bisimulation up to store, $\mathcal{C}$ and expansion. More details can be found in Appendix E.

Second proof. Here we also preprocess the terms using Lemma 6, to add a few artificial internal steps to N, so that we can carry out the reminder of the proof using strong bisimilarity, which enjoys more up-to techniques than weak bisimilarity:

– $M \approx M' \triangleq \lambda g.\nu \ell\ \ell := 0; g(\mathsf{incr}_\ell); \mathsf{test}_\ell I$,
– $N \approx N'' \triangleq \lambda g.I; I; g(\mathsf{incr}_0); \mathsf{test}_0 I$.

where incr_0 and test_0 just return I on any input, taking the same number of internal steps as incr_ℓ and test_ℓ. We show that $M' \sim N''$ by proving that the following relation $\mathcal{R}$ is a strong bisimulation *up to unfolding, store, weakening, strengthening, transitivity and context* (a technique unsound in the weak case):

$$\mathcal{S} \triangleq \{(M', N'')\} \cup \{(\ell \mapsto 2n, \mathsf{incr}_\ell, \mathsf{test}_\ell)\,, (\emptyset, \mathsf{incr}_0, \mathsf{test}_0) \ \text{ s.t. } n \in \mathbb{N}\}$$

This relation uses a single location; there is one pair for each integer that can be stored in the location. In the diagram-chasing arguments for $\mathcal{S}$, essentially a pair of derivatives is proved to be related under the function

$$\mathbf{sp} \circ \mathbf{sp} \circ \mathsf{star} \circ (\mathsf{str} \cup \mathsf{store} \cup \mathcal{C} \cup \mathsf{w})^\omega$$

where $\mathsf{star} : \mathcal{R} \mapsto \mathcal{R}^\star$ is the reflexive-transitive closure function. (Again, we refer to Appendix E for more details.)

The difference between the relation $\mathcal{R}$ in the first proof and the proofs in [10, 21] is that $\mathcal{R}$ only requires locations that appear free in the tested terms; in contrast, the relations in [10, 21] need to be closed under all possible extensions of the store, including extensions in which related locations are mapped onto arbitrary context-closures of related values. We avoid this thanks to the up-to store function. The reason why, both in [10, 21] and in the first proof above, several locations have to be considered is that, with bisimulations akin to environmental bisimulation, the input for a function is built using the values that occur in the candidate relation. In our example, this means that the input for a function can be a context-closure of M and N; hence uses of the input may cause several evaluations of M and N, each of which generates a new location. In this respect, it is surprising that our second proof avoids multiple allocations (the candidate relation $\mathcal{S}$ only mentions one location). This is due to the massive combination of up-to techniques whereby, whenever a new location is created, a double application of up to context (the 'double' is obtained from up-to transitivity) together with some administrative work (given by the other techniques) allows us to absorb the location.

References

1. M. Abadi and A.D. Gordon. A bisimulation method for cryptographic protocols. In Chris Hankin, editor, *ESOP'98*, volume 1381 of *LNCS*, pages 12–26. Springer, 1998.
2. S. Abramsky. The lazy lambda calculus. In D. Turner, editor, *Research Topics in Functional Programming*, pages 65–116. Addison-Wesley, 1989.
3. S. Arun-Kumar and M. Hennessy. An efficiency preorder for processes. *Acta Informatica*, 29:737–760, 1992.
4. K. Chaudhuri, M. Cimini, and D. Miller. Formalization of the bisimulation-up-to technique and its meta theory. Draft, 2014.
5. D. Hirschkoff. A full formalisation of pi-calculus theory in the calculus of constructions. In *TPHOLs*, volume 1275 of *LNCS*, pages 153–169. Springer, 1997.
6. C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis. The power of parameterization in coinductive proof. In *POPL*, pages 193–206. ACM, 2013.

7. A. Jeffrey and J. Rathke. Towards a theory of bisimulation for local names. In *LICS*, pages 56–66, 1999.
8. V. Koutavas and M. Hennessy. First-order reasoning for higher-order concurrency. *Computer Languages, Systems & Structures*, 38(3):242–277, 2012.
9. V. Koutavas, P. B. Levy, and E. Sumii. From applicative to environmental bisimulation. *Electr. Notes Theor. Comput. Sci.*, 276:215–235, 2011.
10. V. Koutavas and M. Wand. Small bisimulations for reasoning about higher-order imperative programs. In *POPL'06*, pages 141–152. ACM, 2006.
11. S.B. Lassen. Relational reasoning about contexts. In *Higher-order operational techniques in semantics*, pages 91–135. Cambridge University Press, 1998.
12. S.B. Lassen. *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, Department of Computer Science, University of Aarhus, 1998.
13. S.B. Lassen. Bisimulation in untyped lambda calculus: Böhm trees and bisimulation up to context. *Electr. Notes Theor. Comput. Sci.*, 20:346–374, 1999.
14. M. Lenisa. *Themes in Final Semantics*. Ph.D. thesis, Università di Pisa, 1998.
15. M. Merro and F. Zappa Nardelli. Behavioral theory for mobile ambients. *J. ACM*, 52(6):961–1023, 2005.
16. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
17. J.Å. Pohjola and J. Parrow. Bisimulation up-to techniques for psi-calculi. Draft, 2014.
18. D. Pous and D. Sangiorgi. Enhancements of the bisimulation proof method. In *Advanced Topics in Bisimulation and Coinduction*. Cambridge University Press, 2012.
19. J. Rot, M. Bonsangue, and J. Rutten. Coalgebraic bisimulation-up-to. In *SOFSEM'13*, volume 7741 of *LNCS*, pages 369–381. Springer, 2013.
20. D. Sangiorgi. On the bisimulation proof method. *J. of MSCS*, 8:447–479, 1998.
21. D. Sangiorgi, N. Kobayashi, and E. Sumii. Environmental bisimulations for higher-order languages. *ACM Trans. Program. Lang. Syst.*, 33(1):5, 2011.
22. D. Sangiorgi and D. Walker. *The Pi-Calculus: a theory of mobile processes*. Cambridge University Press, 2001.
23. E. Sumii and B. C. Pierce. A bisimulation for dynamic sealing. *Theor. Comput. Sci.*, 375(1-3):169–192, 2007.
24. E. Sumii and B. C. Pierce. A bisimulation for type abstraction and recursion. *J. ACM*, 54(5), 2007.
25. N.D. Turner. *The polymorphic pi-calculus: Theory and Implementation*. PhD thesis, Department of Computer Science, University of Edinburgh, 1996.

A First-order bisimulation and up-to techniques

Definition 7 (weak bisimilarity). *A relation $\mathcal{R}$ on processes is a* weak bisimulation *if whenever $P\ \mathcal{R}\ Q$:*

- *if $P \xrightarrow{\mu} P'$ then $Q \stackrel{\hat{\mu}}{\Longrightarrow} Q'$ and $P'\ \mathcal{R}\ Q'$;*
- *if $Q \xrightarrow{\mu} Q'$ then $P \stackrel{\hat{\mu}}{\Longrightarrow} P'$ and $P'\ \mathcal{R}\ Q'$.*

We write $\approx$ for the largest weak bisimulation, and call it weak bisimilarity.

Definition 8 (expansion). *A relation $\mathcal{R}$ on processes is an* expansion relation *if whenever $P\ \mathcal{R}\ Q$:*

- *if $P \xrightarrow{\mu} P'$ then $Q \xrightarrow{\hat{\mu}} Q'$ and $P'\ \mathcal{R}\ Q'$.*
- *if $Q \xrightarrow{\mu} Q'$ then $P \stackrel{\mu}{\Longrightarrow} P'$ and $P'\ \mathcal{R}\ Q'$;*

We write $\gtrsim$ for the largest expansion relation, and simply call it expansion.

B The π-calculus

The syntax of the π-calculus is the following:

$$P ::= 0 \mid a(b).P \mid \overline{a}b.P \mid P|P \mid \boldsymbol{\nu} a\, P \mid\ !P$$

(other operators, such as matching and mismatching, could be added). The operational semantics is described by the rules for $\longmapsto_\pi$ below. We assume that α-convertible terms are identified. The grammar of labels is $\mu ::= \tau \mid ab \mid \overline{a}b \mid \overline{a}(b)$.

$$\textsc{out}\ \frac{}{\overline{a}b.P \stackrel{\overline{a}b}{\longmapsto}_\pi P} \qquad \textsc{inp}\ \frac{}{a(b).P \stackrel{ac}{\longmapsto}_\pi P\{c/b\}} \qquad \textsc{sum-l}\ \frac{P \stackrel{\mu}{\longmapsto}_\pi P'}{P + Q \stackrel{\mu}{\longmapsto}_\pi P'}$$

$$\textsc{par-l}\ \frac{P \stackrel{\mu}{\longmapsto}_\pi P'}{P \mid Q \stackrel{\mu}{\longmapsto}_\pi P' \mid Q}\ \mathrm{bn}(\mu) \cap \mathrm{fn}(Q) = \emptyset \qquad \textsc{comm-l}\ \frac{P \stackrel{\overline{a}b}{\longmapsto}_\pi P' \qquad Q \stackrel{ab}{\longmapsto}_\pi Q'}{P \mid Q \stackrel{\tau}{\longmapsto}_\pi P' \mid Q'}$$

$$\textsc{close-l}\ \frac{P \stackrel{\overline{a}(b)}{\longmapsto}_\pi P' \qquad Q \stackrel{ab}{\longmapsto}_\pi Q'}{P \mid Q \stackrel{\tau}{\longmapsto}_\pi \boldsymbol{\nu} b\,(P' \mid Q')}\ b \notin \mathrm{fn}(Q) \qquad \textsc{res}\ \frac{P \stackrel{\mu}{\longmapsto}_\pi P'}{\boldsymbol{\nu} a\, P \stackrel{\mu}{\longmapsto}_\pi \boldsymbol{\nu} a\, P'}\ a \notin \mathrm{n}(\mu)$$

$$\textsc{open}\ \frac{P \stackrel{\overline{a}b}{\longmapsto}_\pi P'}{\boldsymbol{\nu} b\, P \stackrel{\overline{a}(b)}{\longmapsto}_\pi P'}\ a \neq b \qquad \textsc{rep}\ \frac{P \mid\ !P \stackrel{\mu}{\longmapsto}_\pi P'}{!P \stackrel{\mu}{\longmapsto}_\pi P'}$$

Definition 9 (Bisimilarity in π). *A relation $\mathcal{R}$ is a strong early bisimulation in the π-calculus if, whenever $P\ \mathcal{R}\ Q$:*

1. *if $P \stackrel{\overline{a}(b)}{\longmapsto}_\pi P'$ and $b \notin \mathrm{fn}(Q)$ then $Q \stackrel{\overline{a}(b)}{\longmapsto}_\pi Q'$ for some Q' such that $P'\ \mathcal{R}\ Q'$,*
2. *if $P \stackrel{\mu}{\longmapsto}_\pi P'$ and μ is not a bound output, then $Q \stackrel{\mu}{\longmapsto}_\pi Q'$ for some Q' such that $P'\ \mathcal{R}\ Q'$,*

3. the converse of (1) and (2), on Q.

Early bisimilarity, $\sim^{\mathrm{e}}$, is the union of all early bisimulations.

The weak version of early bisimilarity, *weak early bisimilarity*, written $\approx^{\mathrm{e}}$, is obtained in the standard way: the transition $Q \stackrel{\overline{a}(b)}{\longmapsto}_\pi Q'$ in clause (1) is replaced by $Q \stackrel{\overline{a}(b)}{\Longmapsto}_\pi Q'$; and similarly the transition $Q \stackrel{\mu}{\longmapsto}_\pi Q'$ in (2) is replaced by $Q \stackrel{\hat{\mu}}{\Longrightarrow} Q'$.

C The λ-calculus

The set Λ of pure λ-terms is defined by:

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

We write Λ^0 for the subset of closed terms. The *call-by-name reduction relation* $\longmapsto_{\mathsf{n}}$ is the least relation over Λ^0 that is closed under the following rules.

$$\frac{}{(\lambda x.M)N \longmapsto_{\mathsf{n}} M\{^N\!/\!_x\}} \qquad \frac{M \longmapsto_{\mathsf{n}} M'}{MN \longmapsto_{\mathsf{n}} M'N}$$

We write $\Longmapsto_{\mathsf{n}}$ for the reflexive and transitive closure of $\longmapsto_{\mathsf{n}}$. The values are the terms of the form $\lambda x.M$. In call-by-name *evaluation contexts* are described by the following grammar:

$$C := CM \mid [\cdot]$$

(Symbol C is used also for arbitrary contexts; it will be explicitly indicated when C refers to evaluation contexts.)

D Imperative call-by-value λ-calculus

Here we give more details on a few results for the imperative λ-calculus ΛR. Specifically, first the relationship between contextual equivalence in ΛR and bisimilarity on the target first-order LTS (Theorem 6 of the main text); then Lemma 6 of the main text, and then Theorem 7 that is used in examples (as an instantiation of Lemma 6).

Lemma 7. *If $(s; M) \approx (t; N)$ then $(s; M) \equiv (t; N)$.*

Proof. Let E be an evaluation context. Since $\mathcal{C}_{\mathsf{e}}$ is compatible up to (Lemma 5) by Lemma 2 we know $\approx$ is a $\mathcal{C}_{\mathsf{e}}$-congruence, hence $(s; E[M]) \approx (t; E[N])$.

Suppose now $(s; E[M])$ reaches a value. Then we can derive a $\stackrel{i,C,\mathsf{cod}(r)}{\Longrightarrow}$ transition from it, and by weak bisimulation we can derive the same transition from $(t; E[N])$, meaning it also reaches a value. This means $(s; E[M])\Downarrow$ implies $(t; E[N])\Downarrow$ and we conclude by symmetry.

Lemma 8. *If $(s; M) \equiv (t; N)$ then $(s; M) \approx (t; N)$.*

Proof. We prove a slightly stronger result, that the relation $\mathcal{R}$ below is a weak bisimulation (E ranges over location-free evaluation contexts).

$$\mathcal{R} \triangleq \left\{((s;\Gamma,M),(t;\Delta,N)) \text{ s.t. } \forall E\ (s;E[M,\Gamma])\Downarrow \text{ iff } (t;E[N,\Delta])\Downarrow\right\} \qquad (4)$$

Case 1: τ transition. When $(s;M) \longmapsto_{\mathsf{R}} (s';M)$ we can easily see that $(s;E[M,\Gamma])\Downarrow$ iff $(s';E[M',\Gamma])\Downarrow$ by determinism of $\longmapsto_{\mathsf{R}}$, so we have in fact $(s';\Gamma,M')\ \mathcal{R}\ (t;\Delta,N)$ and the transition $\xrightarrow{\tau}$ is caught up with no transition at all.

Case 2: $i,C,\mathtt{cod}(r)$ transition from $(s;\Gamma,M)$ to $(s';\Gamma'',M')$. (verifying $(*)$ below) means that M is a value V and thus $(s;M)\Downarrow$. By choosing $E=[\cdot]_1$ and $u=\emptyset$ in (4) we know that $(t;N)\Downarrow$ and thus $(t;N)\Longmapsto_{\mathsf{R}} (t';W)$ for some value W.

Hence, we easily get the weak transition $(t;\Delta,N) \overset{i,C,\mathtt{cod}(r)}{\Longrightarrow} (t'';\Delta'',N')$ through $(t';\Delta,W)$, verifying $(**)$ below.

$$\Gamma' = \Gamma, V \quad \Gamma'' = \Gamma', \mathsf{getset}(r) \quad (r[\Gamma''] \uplus s; \Gamma'_i(C[\Gamma'])) \longmapsto_{\mathsf{R}} (s';M') \qquad (*)$$
$$\Delta' = \Delta, W \quad \Delta'' = \Delta', \mathsf{getset}(r) \quad (r[\Delta''] \uplus t'; \Delta'_i(C[\Delta''])) \longmapsto_{\mathsf{R}} (t'';N') \qquad (**)$$

We prove now $(s';\Gamma'',M')\ \mathcal{R}\ (t'';\Delta'',N')$. Let E be any location-free evaluation context, we prove:

$$(s';E[M',\Gamma'])\Downarrow \quad \text{iff} \quad (t'';E[N',\Delta'])\Downarrow \quad . \qquad (5)$$

Backtracking one step using $(*)$ and $(**)$ it is enough to prove

$$(r[\Gamma''] \uplus s; E[\Gamma'_i(C[\Gamma'']),\Gamma''])\Downarrow \quad \text{iff} \quad (r[\Gamma''] \uplus t'; E[\Delta'_i(C[\Delta'']),\Delta''])\Downarrow \qquad (6)$$

which we do by exhibiting an evaluation context F for which we already have (by instantiating with F the definition of $\mathcal{R}$) the equation (7) below, and that each member of (6) is a derivative of the corresponding member of (7).

$$(s;F[M,\Gamma])\Downarrow \quad \text{iff} \quad (t;F[N,\Delta])\Downarrow \qquad (7)$$

For that we choose

$$F \triangleq \ \mathtt{let}\ x = [\cdot]_1\ \mathtt{in}\ \nu\ell_1 \ldots \nu\ell_n\ \ell_1 := C_1^\bullet;\ \ \ldots; \ell_n := C_n^\bullet;\ E^\bullet[[\cdot]_i^\bullet(C^\bullet), -]$$

where r is the collection of $\ell_i \mapsto C_i$ and we write $D^\bullet$ for a context D where the holes destined to get_{ℓ_i} and set_{ℓ_i} are already filled, and the holes destined to get the values V and W are filled with x.

The lemma below is Lemma 6 of the main text.

Lemma 9. *Let $\asymp$ be any of the relations $\sim, \approx$, and $\gtrsim$. Suppose L, R are ΛR terms with $(s;\Gamma,L) \asymp (s;\Gamma,R)$ for all environments Γ and stores s. Then also $(s;\Gamma,C[L]) \asymp (s;\Gamma,C[R])$, for any store s, environment Γ and context C.*

We give the proof for $\gtrsim$ as it is the most general case. Also remark that the last Γ is not necessary, as it can be encoded into the C.

The proof goes as follows: (1) we first prove a simplified result in which the context C is an evaluation context, using techniques $\mathcal{C}_{\mathsf{e}}$ and store. (2) We then exploit (1) to derive another partial result where C is a context whose holes are *not* in evaluation position, and achieve the proof using up to expansion and (1) when a hole is freed.

These classes of contexts are enough to cover all cases and the proofs (1) and (2) focus on very different parts of the problem. This separation is necessary: side effects of the store would make quite convoluted a naive bisimulation candidate, on which case analyses prove difficult.

Lemma 10 handles (1) and Lemma 11 handles (2).

Lemma 10. *Suppose that for all s and Γ, we have $(s;\Gamma,L)\gtrsim(s;\Gamma,R)$. Then for all s, Γ and evaluation context F with free locations, $(s;\Gamma,F[L])\gtrsim(s;\Gamma,F[R])$.*

Proof. Let A be the list of set_ℓ and get_ℓ for all location ℓ in F. Then it is easy to get F' from F such that $F=F'[-,A]$ and F' is location-free. By hypothesis we know $(s;\Gamma,A,L)\gtrsim(s;\Gamma,A,R)$ on which we apply precongruence for evaluation contexts $\mathcal{C}_{\mathsf{e}}$ to get $(s;\Gamma,A,F'[L,A])\gtrsim(s;\Gamma,A,F'[R,A])$. By weakening w we get $(s;\Gamma,F'[L,A])\gtrsim(s;\Gamma,F'[R,A])$.

Lemma 11. *Let L, R be ΛR terms with $(s;\Gamma,L)\gtrsim(s;\Gamma,R)$ for all environment Γ and store s. Then Suppose C is a multi hole context, such that no hole is in evaluation position in C. Then for all store s we know $(s;C[L])\gtrsim(s;C[R])$.*

Proof. Let $\mathcal{R}$ relate each configuration $((\ell\mapsto C_v^\ell[L])_\ell;\tilde{C}_v[L],C[L])$ to the one where R replaces L: $((\ell\mapsto C_v^\ell[R])_\ell;\tilde{C}_v[R],C[R])$ for all $(C_v^\ell)_\ell$ and $\tilde{C}_v$ families of value contexts (of the form $\lambda x.C'$), and C context without any hole in evaluation position. For short we write s_L, s_R, Γ_L, and Γ_R the corresponding stores and environments. We run simultaneously the transitions from both sides $(s_L;\Gamma_L,C[L])$ and $(s_R;\Gamma_R,C[R])$ as they have always the same shape.

We prove $\mathcal{R}$ is an expansion relation up to expansion. We rely on the fact that L and R will never be run directly in this proof.

Case 1: τ action. Since in L in $C[L]$ (and R in $C[R]$) is not in evaluation position, both sides will do the same kind of transition, being completely oblivious to L/R. The resulting configurations will be $(s'_L;\Gamma_L,C'[L])$ and $(s'_R;\Gamma_R,C'[R])$. (Even if a set_ℓ or a get_ℓ is involved, L/R part may go to or from the store, but this will keep the same shape.)

The only part of the invariant of the relation that is not maintained is that L/R may appear in evaluation position, if $C'[L]=E[L,L]$ (where $[\cdot]_1$ is in evaluation position and $[\cdot]_2$ may appear everywhere).

In this case, we remark that $F\triangleq E[-,L]$ is an evaluation context, on which we can apply Lemma 10 to have $(s'_L;\Gamma_L,E[L,L])\gtrsim(s'_L;\Gamma_L,E[R,L])$ and now since R is not in evaluation position, the context $C_2=E[R,-]$ is a context with no hole in evaluation position, hence $(s'_L;\Gamma_L,E[R,L])\ \mathcal{R}\ (s'_R;\Gamma_R,E[R,R])$ and we have closed the diagram.

Note that it may happen that even if $E[L,-]$ doesn't have holes in evaluation position, $E[R,-]$ does. In this case, we just use Lemma 10 while there are still such holes, and the progression to $\gtrsim\mathcal{R}$ still holds ($\gtrsim$ is transitive).

Case 2: visible action. First, $C[L]$ is a value iff $C[R]$ is a value so they have the same visible actions of the form $i, D, \mathtt{cod}(r)$. We end up in the same shape of configurations we had for the τ transition above, and proceed the same to close the diagram.

Finally we have proven that $\mathcal{R}$ progresses to $\gtrsim\mathcal{R}$ (expansion up to expansion). In the strong case, we would have proven $\mathcal{R}$ progresses to $\sim\mathcal{R}$, and in the weak case we would have proven that it progresses to both $\approx\mathcal{R}$ *and* $\mathcal{R}\approx$, which is necessary because in the weak case, one can use "up to $\approx$" only when $\approx$ is not on the same side as the challenge.

In the following $\mathcal{R}^+$ is the transitive closure of $\mathcal{R}$. We prove here Theorem 7 (using Lemma 12) since we use (simple instances of) it in Section 6.

Lemma 12. *Suppose that E and E' are evaluation contexts and that for all value V value and store s, we have $(s; E[V]) \longmapsto_{\mathsf{R}}^{+} (s; E'[V])$. Then for all environment Γ and store s, we have $(s; \Gamma, E[M]) \gtrsim (s; \Gamma, E'[M])$.*

Proof. For a given Γ we consider $\mathcal{R} = \{(s; \Gamma, E[M]), (s; \Gamma, E'[M]) \mid s, M\}$ and the transitions from both sides:

1. when M is not a value, $(s; M) \longmapsto_{\mathsf{R}} (s'; M')$ and the only transition from both sides is a τ staying knowingly in the relation.
2. (left to right) when $M = V$ by hypothesis $(s; \Gamma, E[V]) \xrightarrow{\tau}{}^{+} s(; \Gamma, E'[V])$ so the first transition from the left-hand side is a τ. We use up to expansion to reach $(s; \Gamma, E'[V])$ which is equal to the right-hand side, and conclude up to reflexivity.
3. (right to left) when $M = V$ and the right-hand side makes some transition $(s; \Gamma, E'[V]) \xrightarrow{\alpha} (s'; \Gamma', N')$ we know in fact that $(s; \Gamma, E[V]) \xrightarrow{\tau}{}^{+} (s; \Gamma, E'[V])$ so $(s; \Gamma, E'[V]) \Longrightarrow\!\!\xrightarrow{\alpha} (s'; \Gamma', N')$ and we conclude again up to reflexivity.

We proved $\mathcal{R}$ is an expansion relation up to expansion and reflexivity.

Remark 3. To get to Theorem 7 we combine (in Lemma 6) proofs for evaluation contexts (Lemma 12) and for non-evaluation contexts (Lemma 11). This separation is critical, as handling all contexts together would yield a much bigger and error-prone bisimulation candidate as L and R in Lemma 11 would be replaced by all intricate combinations of E and E'.

Remark 4. In the proofs leading to Theorem 7 we universally quantify over contexts several times, but we use up to context techniques only a few times. This makes sense, as those are arbitrary contexts with locations containing arbitrary terms that are not necessarily values; we needed tight control over them, and the resulting fine-tuned proof can now be used as a black box.

Remark 5. To see why separating the proof into Lemma 10 Lemma 11 is necessary one must go through several naive steps when expanding the candidate relation relating $(\emptyset; C[(\lambda x.E[x])M])$ to $(\emptyset; C[E[M]])$.

– there can be a location ℓ both in C and M. For instance, C could be $\mathsf{set}_\ell(C_v); C_2$ where C_v is a value context, so the store must be able to contain $s = (\ell \mapsto C_v[(\lambda x.E[x])M])$ on the left, where it contains $s' = (\ell \mapsto C_v[E[M]])$ on the right.
– then C can be $[\cdot]$ so we must be able to compare $(s; (\lambda x.E[x])M)$ to $(s'; E[M])$ which calls on how to relate $(s; M)$ to $(s'; M)$, for instance it implies proving that either both or none reach a value, which we don't know yet, because that is similar to what we already intended to prove (we get into a circular argument).

Theorem 7. *Suppose that E and E' are evaluation contexts and that for all value V value and store s, we have $(s; E[V]) \longmapsto_{\mathsf{R}}^{*} (s; E'[V])$. Then for all environment s and context C, we have $(s; C[E[M]]) \gtrsim (s; C[E'[M]])$.*

Proof. Consequence of Lemma 12 and Lemma 6.

E Example from Section 6

Continuing from Section 6, we show that the relation

$$\mathcal{R} = \left\{ \left(s; M', (\mathsf{incr}_\ell, \mathsf{test}_\ell)_{\ell\in\tilde{\ell}}\right), \left(\emptyset; N', (F, F)_{\ell\in\tilde{\ell}}\right) \text{ s.t. } \forall \ell \in \tilde{\ell}, \quad s(\ell) \text{ is even} \right\}$$

is a weak bisimulation up to store, $\mathcal{C}$ and expansion. We write $(s; \Gamma_{\tilde{\ell}})$ for the left-hand side of a pair in $\mathcal{R}$ and $(\emptyset; \Delta_{\tilde{\ell}})$ for the right-hand side.

Consider a transition $1, C, \mathtt{cod}(r)$ from M' and N'. We write below Γ' for $\Gamma_{\tilde{\ell}}, \mathsf{getset}(r)$ and Δ' for $\Delta_{\tilde{\ell}}, \mathsf{getset}(r)$.

– $(s; \Gamma_{\tilde{\ell}}) \xrightarrow{1,C,\mathtt{cod}(r)} (s \uplus r[\Gamma']; \Gamma', \nu\ell\ \ell := 0; C[\Gamma'](\mathsf{incr}_\ell); \mathsf{test}_\ell I)$
– $(\emptyset; \Delta_{\tilde{\ell}}) \xrightarrow{1,C,\mathtt{cod}(r)} (r[\Delta']; \Delta', C[\Delta'](F); FI)$

In the first line, we make the configuration run two τ transitions, so that $\nu\ell$ and $\ell := 0$ get executed. Now we have a new store $s' = s \uplus (\ell \mapsto 0)$ ($s'(\ell)$ is even, so in this respect we stay in the bisimulation candidate).

Now the main term is $C[\Gamma'](\mathsf{incr}_\ell); \mathsf{test}_\ell I)$ which can be rewritten to $D[\Gamma_{\tilde{\ell}\cdot\ell}, \mathsf{getset}(r)]$ for some context D. On the right-hand side $C[\Delta'](F); FI$ can be rewritten to $D[\Delta_{\tilde{\ell}\cdot\ell}, \mathsf{getset}(r)]$ as well. By construction $(s'; \Gamma_{\tilde{\ell}\cdot\ell})\ \mathcal{R}\ (\emptyset; \Delta_{\tilde{\ell}\cdot\ell})$ hence

$$(s' \uplus r[\Gamma']; \Gamma_{\tilde{\ell}\cdot\ell}, \mathsf{getset}(r)) \quad \mathsf{store}(\mathcal{R}) \quad (r[\Delta']; \Delta_{\tilde{\ell}\cdot\ell}, \mathsf{getset}(r)) .$$

Now we apply $\mathcal{C}$ with the context D, then the weakening w to remove incr_ℓ and test_ℓ to reach the pair we wanted.

Now that we handled M' and N', let us look at any transition $i, C, \mathtt{cod}(r)$ coming from some incr_ℓ (and F on the other side). It will result in I on both sides (the argument C is discarded), with $s(\ell)$ being updated to $s(\ell) + 2$ which stays in the relation. The store is augmented with $r[\Gamma']$ and $r[\Delta']$ and the environment with $\mathsf{getset}(r)$ which can be safely removed by the store technique as we did before. The same is done when a test_ℓ is run: both sides reduce to I, the argument is discarded, and the r part of the transition is garbage-collected.

We now present some details for the second proof of the example. We show that the relation

$$\mathcal{S} = \{(M', N'')\} \cup \{(\ell \mapsto 2n; \mathsf{incr}_\ell, \mathsf{test}_\ell)\,, (\emptyset; \mathsf{incr}_0, \mathsf{test}_0) \;\text{ s.t. }\; n \in \mathbb{N}\}$$

is a strong bisimulation *up to unfolding, store, weakening, strengthening, transitivity and context.*

This up-to technique, unsound in the weak case (transitivity is unsound), is powerful enough to make the bisimulation considerably smaller. Proving that the second member of $\mathcal{S}$ progresses to itself (up to store) is straightforward. We focus on the following transitions from M' and N'':

$$\begin{aligned}(\emptyset, M') &\xrightarrow{1,C,\mathtt{cod}(r)} (\ r[\Gamma];\ \ \Gamma, \nu\ell\ \ell := 0; C[\Gamma](\mathsf{incr}_\ell); \mathsf{test}_\ell I\ ) \triangleq H_1\\ (\emptyset, N'') &\xrightarrow{1,C,\mathtt{cod}(r)} (\ r[\Delta];\ \ \Delta, I; I; \qquad C[\Delta](\mathsf{incr}_0); \mathsf{test}_0 I\ ) \triangleq H_2\end{aligned}$$

where $\Gamma = M', \mathsf{getset}(r)$ and $\Delta = N'', \mathsf{getset}(r)$. We use $\mathbf{sp}$ as an up-to technique[5] twice to run two steps of reduction on both sides:

$$H_1 \xrightarrow{\tau}\xrightarrow{\tau} H_1' \qquad \text{and} \qquad H_2 \xrightarrow{\tau}\xrightarrow{\tau} H_2'\,.$$

This way we trigger $\nu\ell$ and $\ell := 0$ and obtain two configurations H_1' and H_2' that can be related using a few up-to functions:

$$\begin{aligned}&(\ r[\Gamma] \uplus (\ell \mapsto 0);\ \ \Gamma,\, C[\Gamma](\mathsf{incr}_\ell); \mathsf{test}_\ell I\ ) = H_1'\\ \mathsf{w}(\mathcal{C}(\mathsf{store}(\mathsf{str}(\mathcal{S})))) \quad &(\ r[\Gamma];\ \ \Gamma,\, C[\Gamma](\mathsf{incr}_0); \mathsf{test}_0 I\ )\\ \mathcal{C}(\mathsf{store}(\mathcal{S})) \quad &(\ r[\Delta];\ \ \Delta,\, C[\Delta](\mathsf{incr}_0); \mathsf{test}_0 I\ ) = H_2'\ .\end{aligned}$$

We detail below how we go from the first to the second line. We write $\Gamma_\ell \triangleq \mathsf{incr}_\ell, \mathsf{test}_\ell$ and $\Gamma_0 \triangleq \mathsf{incr}_0, \mathsf{test}_0$.

$$\begin{array}{lcl}
(\ell \mapsto 0; \Gamma_\ell) & \mathcal{S} & (\emptyset; \Gamma_0)\\
(\ell \mapsto 0; \Gamma_\ell, M') & \mathsf{str}(-) & (\emptyset; \Gamma_0, M')\\
(r[\Gamma] \uplus \ell \mapsto 0; \Gamma_\ell, \Gamma) & \mathsf{store}(-) & (r[\Gamma]; \Gamma_0, \Gamma)\\
(r[\Gamma] \uplus \ell \mapsto 0; \Gamma_\ell, \Gamma, C[\Gamma](\mathsf{incr}_\ell); \mathsf{test}_\ell I)) & \mathcal{C}(-) & (r[\Gamma]; \Gamma_0, \Gamma, C[\Gamma](\mathsf{incr}_0); \mathsf{test}_0 I)\\
(r[\Gamma] \uplus \ell \mapsto 0; \quad \Gamma, C[\Gamma](\mathsf{incr}_\ell); \mathsf{test}_\ell I)) & \mathsf{w}(-) & (r[\Gamma]; \quad \Gamma, C[\Gamma](\mathsf{incr}_0); \mathsf{test}_0 I)
\end{array}$$

Going from the second to the third line is easier:

$$\begin{array}{lcl}
(\emptyset; M') & \mathcal{S} & (\emptyset; N'')\\
(r[\Gamma]; \Gamma) & \mathsf{store}(\mathcal{S}) & (r[\Delta]; \Delta)\\
(r[\Gamma]; \Gamma, C[\Gamma](\mathsf{incr}_0); \mathsf{test}_0 I) & \mathcal{C}(\mathsf{store}(\mathcal{S})) & (r[\Delta]; \Delta, C[\Delta](\mathsf{incr}_0); \mathsf{test}_0 I)
\end{array}$$

Finally we proved that $H_1\ f(\mathcal{S})\ H_2$ where $f = \mathbf{sp} \circ \mathbf{sp} \circ \mathsf{star} \circ (\mathsf{str} \cup \mathsf{store} \cup \mathcal{C} \cup \mathsf{w})^\omega$ is a compatible function, and hence $\mathcal{S} \rightsquigarrow_{\mathbf{sp}} f(\mathcal{S}) \cup \mathsf{store}(\mathcal{S})$ (not forgetting the second member of $\mathcal{S}$).

To conclude, $\mathcal{S}$, as a strong bisimulation up to (unfolding, store, weakening, strengthening, transitivity and context), is included in $\sim$.

[5] If $\xrightarrow{\tau}$ is deterministic then $(\xrightarrow{\tau} \mathcal{R} \xleftarrow{\tau}) \subseteq \mathbf{sp}(\mathcal{R})$.