

On sequentiality and well-bracketing in the π -calculus[★]Daniel Hirschhoff^a, Enguerrand Prebet^{a,b}, Davide Sangiorgi^{id c,d,*}^a ENS de Lyon, France^b Karlsruhe Institute of Technology, Germany^c Università di Bologna, Italy^d INRIA Sophia Antipolis, France

ARTICLE INFO

Keywords:

Behavioural equivalence

Bisimilarity

Process calculi

Sequentiality

Well-bracketing

ABSTRACT

The π -calculus is used as a model for programming languages. Its contexts exhibit arbitrary concurrency, making them very discriminating. This may prevent validating desirable behavioural equivalences in cases when more disciplined contexts are expected.

In this paper we focus on two such common disciplines: sequentiality, meaning that at any time there is a single thread of computation, and well-bracketing, meaning that calls to external services obey a stack-like discipline. We formalise the disciplines by means of type systems. The main focus of the paper is on studying the consequence of the disciplines on behavioural equivalence. We define and study labelled bisimilarities for sequentiality and well-bracketing. These relations are coarser than ordinary bisimilarity. We prove that they are sound for the respective (contextual) barbed equivalence, and also complete under a certain technical condition.

We show the usefulness of our techniques on a number of examples, that have mainly to do with the representation of functions and store.

1. Introduction

The π -calculus has been advocated as a model to give semantics to, and reason about, various forms of programming languages, including those with higher-order features. Strengths of the π -calculus are its rich algebraic theory and its wide spectrum of proof techniques. Concurrency is at the heart of the π -calculus: computation is interaction between concurrent processes. The operators of the calculus are simple (parallelism, input, output, restriction being the main ones) and unconstrained. This yields an amazing expressive power — the calculus can model a variety of programming idioms [1]. However, this also makes the contexts of the calculus very discriminating; as a consequence, behavioral equivalences, which are supposed to be preserved by all the contexts of the calculus, are rather demanding relations.

Higher-level languages may be syntactically quite different from a language for pure concurrency such as the π -calculus. For instance, the paradigmatic higher-order programming language, the λ -calculus, is a pure calculus of functions and, in both its call-by-name and call-by-value variants, is sequential — it is even deterministic. A variety of extensions of it have been considered; examples of additional features are references, control operators, non-determinism, (constrained) forms of concurrency. The specific set of syntactic features chosen for the language determines the ways in which the contexts of the language may interact with the terms. In any case, the patterns of interaction are usually more disciplined than those that arise in π -calculus representations of those

[★] This paper is a revised and extended version of a paper appeared in the proc. LICS '21, 1–13, ACM/IEEE.

* Corresponding author.

E-mail address: davide.sangiorgi@unibo.it (D. Sangiorgi).

terms. Thus there are λ -terms that are indistinguishable within the (pure) λ -calculus whose π -calculus images can be separated by appropriate π -contexts.

A well-known way of imposing a discipline to the π -calculus is to equip it with a type system. Such systems are intended to capture communication patterns that occur frequently when programming in the π -calculus. A number of type systems have been considered: e.g., capability types (formalizing the intended I/O usage of names that are exchanged among processes), linearity (formalizing the property that certain names may be used at most once), session types (formalizing the communication protocols in the dialogs between two or more processes), and so on [2–6]. Type systems have also been designed to capture specific properties of processes, such as termination, deadlock-freedom, lock-freedom [7–11]. Types impose constraints on the set of legal contexts in which well-typed terms are supposed to be used; this can make behavioral equivalences usefully coarser.

A further step is then to tune the proof techniques of the π -calculus to such type systems, so to be able to actually prove the behavioral equalities that only hold in the presence of types. Typically this is investigated in the coinductive setting of bisimilarity, and achieved by refining and/or modifying the standard bisimilarity clauses so to take the usage of types into account. The resulting bisimilarity should be sound with respect to contextually-defined forms of bisimilarity such as *barbed equivalence* (or *congruence*); ideally, it should also be complete.

In barbed equivalence, the bisimulation game is played only on internal actions, and certain success signals, the barbs, are used to monitor the computation. In the standard barbed equivalence, an arbitrary context may be added, once (at the beginning), on top of the tested processes. In *reduction-closed* barbed equivalence [1,12], the context may be dynamically updated, by adding further components during the computation. Reduction-closed barbed equivalence usually allows simpler proofs of completeness, and does not require any hypothesis of image-finiteness on the state space of the tested processes. In contrast, standard barbed equivalence is more robust — reduction-closed barbed equivalence may sometimes be over-discriminating [13].

In this paper we focus on the π -calculus representation of *sequentiality* and *well-bracketing*. ‘Sequentiality’ intuitively indicates the existence of a single thread of computation. ‘Well-bracketing’ is a terminology borrowed from game semantics, and used to refer to a language without control operators, in which the call-return interaction behavior between a term and its context follows a stack discipline. Our main objectives are to define bisimilarity-based proof techniques for type systems in the π -calculus that formalize the sequentiality and well-bracketing notions. We actually work with the *asynchronous* π -calculus, $A\pi$, as this is the calculus that is usually adopted in the literature for modeling higher-order languages.

In $A\pi$, sequentiality is the property that, at any time, at most one process is *active*, or *carries the thread*; that is, the process has the control on the computation and decides what the next computation step can be. In other words, we never find two sub-components of a system both of which contain an *interaction redex* (a pair of an input and an output processes at the same name).

In the (standard) encodings of the λ -calculus [14,15], a process is active, i.e., it carries the thread, when it contains an unguarded output particle. Indeed, the π -calculus terms obtained from the encodings give rise to computations in which, syntactically, at any time there is at most one unguarded output particle. An input process that consumes that output will in turn become active.

Our type system is more general, in that we allow also input processes to carry the thread. The type system specifies whether a name may carry the thread in output or in input; we call these names *output-controlled* and *input-controlled*. While the output-controlled are the most important ones (for instance, they play a central role in the modeling of functions), input-controlled names may be useful too, for instance, in the representation of references or locks. A reference ℓ that contains the value n is represented in $A\pi$ by an output particle $\bar{\ell}(n)$; and a process accessing the reference will do so by performing an input at ℓ . Thus an input at ℓ indicates ownership of the current computation thread.

As remarked above, sequentiality implies absence of parallel computation threads. Sequentiality however does not exclude non-determinism. An output particle $\bar{a}(b)$ that owns the thread may have the possibility of interacting with different input processes at a (and symmetrically for input processes owning the thread). Indeed we also admit internal non-determinism (i.e, processes such as $\tau.P + \tau.Q$ that may choose to reduce either to P or to Q without interactions with the environment), both in active and in inactive processes.

The type system for well-bracketing is a refinement of that for sequentiality. Intuitively, well-bracketing means that interactions, in the form of call-return pairs, are properly nested. Imagine a process interrogating services in its environment, where each interrogation is supposed to terminate with an answer. In a well-bracketed system the most recently opened service call must be answered first, akin to a discipline of matching parentheses, or a call stack. A well-bracketed environment cannot return to an older pending call while a newer one is still unanswered. In the type system for well-bracketing, a stack of *continuation names* is used to keep track of the structure of calls and returns among the processes. These stacks are similar to those used in the implementation of compilers for languages (or fragments of languages) adopting well-bracketing, or used in well-bracketed forms of game semantics.

Both sequentiality and well-bracketing put constraints on the possible interactions that a ‘legal’ observer may offer to a tested process. For instance, in the case of sequentiality, if the process has the thread then process and observer may only interact along a name carrying the thread — any other interaction is simply forbidden. Similarly, in the case of well-bracketing, the observer may not offer an answer to an older call while a newer one is still unanswered. As a consequence, sequentiality and well-bracketing have an impact on behavioral equivalences: as the set of legal observers is smaller, the set of pairs of indistinguishable processes is larger.

Finding proof techniques to reason about sequentiality and well-bracketing presents a number of caveats, that have mainly to do with the soundness and completeness of the resulting bisimilarity with respect to barbed equivalence. We briefly discuss below a couple of issues concerning completeness.

In the proof of completeness one has to show that the contexts of the language are at least as discriminating as the labeled bisimilarity. In standard proofs, one defines special contexts that interact with the tested processes and, at the same time, emit certain signals to the outside so to provide information on the kind of interactions that have occurred with the processes. Such behavior of

$$\begin{aligned}
P \mid \mathbf{0} &\equiv P & P \mid Q &\equiv Q \mid P & P \mid (Q \mid R) &\equiv (P \mid Q) \mid R & G + \mathbf{0} &\equiv G & G + G' &\equiv G' + G \\
G + (G' + G'') &\equiv (G + G') + G'' & !a(\tilde{b}).P &\equiv a(\tilde{b}).P \mid !a(\tilde{b}).P & P \mid (\nu a)Q &\equiv (\nu a)(P \mid Q) \text{ if } a \notin \text{fn}(P) \\
(\nu a)(\nu b)P &\equiv (\nu b)(\nu a)P & (\nu a)\mathbf{0} &\equiv \mathbf{0} & [a = a]G &\equiv G
\end{aligned}$$

Fig. 1. Structural congruence in $\mathcal{A}\pi$.

the testing contexts is however inherently concurrent — the context has to interact with the tested processes and, at the same time, emit signals to the outside — and is therefore liable to break the typing discipline for sequentiality (and hence also well-bracketing).

Further problems arise in proofs about reduction-closed barbed equivalence. The reason why completeness proofs for reduction-closed barbed equivalence may be simpler than with standard barbed equivalence is that the testing context may be incrementally adjusted, after every interaction step with the tested processes. This however requires the existence of special components in the contexts to handle the fresh names generated by the tested processes. Specifically, the task of these components is to ensure that new pieces of contexts, added later, will be able to access such fresh names. Again, these components represent parallel threads, and break the sequentiality and well-bracketing disciplines. For this reason in the paper we cannot appeal to reduction-closed forms of barbed equivalence, remaining within the standard notions and therefore requiring an image-finiteness condition.

In the case of well-bracketing the problems above are enhanced by the presence of continuation names. These names are *linear* [3] (they may only be used once), *input receptive* [16] (the input-end of the name should always be available), and output-controlled. This places further constraints on the use of such names within contexts that test the processes. For the above reasons, the completeness proofs for sequentiality and well-bracketing present significant technical differences, both between them and from completeness proofs in the literature.

In the paper we propose labeled bisimilarities that allow us to reason about processes following the sequentiality or well-bracketing disciplines. We prove that the bisimilarities are sound with respect to barbed equivalence. We also establish completeness, on processes with only output-controlled names. We do not know whether completeness holds in the general case, with also input-controlled names. We also study some refinements of the bisimilarities: one is obtained by injecting ideas from bisimilarities for calculi with references [17]; other refinements are forms of ‘up-to techniques’. We illustrate applications of our techniques on a number of examples, most of which have to do with the representation of functions and references. Usually the examples are about equalities that only hold under the sequentiality or well-bracketing disciplines; other examples show that sequentiality and well-bracketing may make equalities simpler to prove because there are fewer observables to take into account.

Paper outline. We introduce some background in Section 2. We study sequentiality in Section 3, and well-bracketing in Section 4: in each case, we present our type system, define an appropriate notion of bisimilarity, and show some examples or laws that we can derive. Related and future works are discussed in Section 5.

2. Background: the (asynchronous) π -calculus

We recall here the standard syntax of the asynchronous π -calculus, $\mathcal{A}\pi$, from [18]:

$$\begin{aligned}
P, Q &::= \bar{a}(\tilde{b}) \mid !a(\tilde{b}).P \mid P \mid Q \mid (\nu a)P \mid G \\
G, G' &::= \mathbf{0} \mid a(\tilde{b}).P \mid \tau.P \mid [a = b]G \mid G + G'
\end{aligned}$$

Names are ranged over by $a, b, \dots$. In prefixes $\bar{a}(\tilde{b})$ and $a(\tilde{b}).P$, name a is the *subject* and $\tilde{b}$ are the *objects*. We use a tilde, like in $\tilde{b}$, for (possibly empty) tuples of names; similarly $(\nu \tilde{a})P$ stands for a sequence of restrictions. As usual, we write $a.P$ and $\bar{a}$ when the object of a prefix is the empty tuple. We use $\sum_{i \in I} G_i$ (resp. $\prod_{i \in I} P_i$) for $G_{i_1} + \dots + G_{i_n}$ (resp. $P_{i_1} \mid \dots \mid P_{i_n}$) where $I = \{i_1, \dots, i_n\}$. We write $P\{a/b\}$ for the result of replacing name b with a in P in a capture-avoiding way. Contexts, C , are processes containing a single occurrence of a special constant, the hole (written $[\cdot]$). The *static* contexts, ranged over by E , have the form $\nu \tilde{a}(P \mid [\cdot])$. In examples, for readability we sometimes use basic data values such as integers and booleans. The definition of structural congruence, written $\equiv$, and of the strong and weak early labeled transitions, written $\xrightarrow{\mu}, \Rightarrow$, and $\xRightarrow{\mu}$, are standard. The definition of structural congruence is reported in Fig. 1.

The (early) Labeled Transition System for $\mathcal{A}\pi$ is reported in Fig. 2, actions being defined by the following grammar:

$$\mu ::= \tau \mid a(\tilde{b}) \mid (\nu \tilde{c})\bar{a}(\tilde{b})$$

We write $\text{fn}(P)$ (resp. $\text{fn}(\mu)$) for the set of free names of P (resp. μ). We have $\text{fn}(a(\tilde{b})) = \{a\} \cup \tilde{b}$, and $\text{fn}((\nu \tilde{c})\bar{a}(\tilde{b})) = (\{a\} \cup \tilde{b}) \setminus \tilde{c}$.

Based on the labeled transitions, we define the following relations: $\Rightarrow$ stands for the reflexive transitive closure of $\xrightarrow{\tau}$; and $\xRightarrow{\mu}$ stands for $\xRightarrow{\mu} \Rightarrow$ (the composition of the three relations). Moreover, $\xRightarrow{\mu}$ is equal to $\xRightarrow{\mu}$ when $\mu \neq \tau$, otherwise it is the same as $\Rightarrow$. We sometimes abbreviate reductions $P \xrightarrow{\tau} P'$ as $P \longrightarrow P'$.

The calculi in the paper will be typed. For simplicity we define our type systems as refinements of the most basic type system for π -calculus, namely Milner’s *sorting* [19], in which names are partitioned into a collection of *types* (or *sorts*), and a sorting function

$$\begin{array}{c}
\text{A}_{\text{INP}} \frac{}{a(\tilde{b}).P \xrightarrow{a(\tilde{c})} P\{\tilde{c}/\tilde{b}\}} \quad \text{A}_{\text{OUT}} \frac{}{\bar{a}(\tilde{b}) \xrightarrow{\bar{a}(\tilde{b})} \mathbf{0}} \quad \text{A}_{\text{OPEN}} \frac{P \xrightarrow{(\nu\tilde{c})\bar{a}(\tilde{b})} P'}{(\nu d)P \xrightarrow{(\nu d, \tilde{c})\bar{a}(\tilde{b})} P'} \text{ if } d \in \text{fn}((\nu\tilde{c})\bar{a}(\tilde{b})) \setminus \{a\} \\
\\
\text{A}_{\text{REP}} \frac{a(\tilde{b}).P \xrightarrow{\mu} P'}{!a(\tilde{b}).P \xrightarrow{\mu} P' \mid !a(\tilde{b}).P} \quad \text{A}_{\text{RES}} \frac{P \xrightarrow{\mu} P'}{(\nu a)P \xrightarrow{\mu} (\nu a)P'} \text{ if } a \notin \text{fn}(\mu) \cup \text{bn}(\mu) \\
\\
\text{A}_{\text{PAR}} \frac{P \xrightarrow{\mu} P'}{P \mid Q \xrightarrow{\mu} P' \mid Q} \text{ if } \text{bn}(\mu) \cap \text{fn}(Q) = \emptyset \quad \text{A}_{\text{COMM}} \frac{P \xrightarrow{a(\tilde{b})} P' \quad Q \xrightarrow{(\nu\tilde{c})\bar{a}(\tilde{b})} Q'}{P \mid Q \xrightarrow{\tau} (\nu\tilde{c})(P' \mid Q')} \text{ if } \tilde{c} \cap \text{fn}(P) = \emptyset \\
\\
\text{A}_{\text{MATCH}} \frac{P \xrightarrow{\mu} P'}{[a = a]P \xrightarrow{\mu} P'} \quad \text{A}_{\text{SUM}} \frac{P \xrightarrow{\mu} P'}{P + Q \xrightarrow{\mu} P'}
\end{array}$$

Fig. 2. Early Labeled Transition Semantics for $\lambda\pi$.

maps name types onto name types. If a name type S is mapped onto a type T , this means that names in S may only carry names in T . We assume that there is a sorting system under which all processes we manipulate are well-typed. We write $\Delta \vdash P$ when process P is well-typed under the typing Δ , and similarly for other objects, such as contexts.

The reference behavioral equivalence for us will be the context-closure of barbed bisimulation. We focus on barbed equivalence (as opposed to barbed congruence) because it is simpler (notably, we do not need to consider issues of closure of the labeled bisimulations under name substitutions). The definition of barbed bisimulation uses the reduction relation $\Rightarrow$ along with an observation predicate $\Downarrow_a$ for each name a , which detects the possibility of performing an output to the external environment along a . Moreover, since we work in a typed setting, such an output should be allowed by the typing of the tested processes. Thus, we write $\Delta \models P \Downarrow_a$ if Δ is a typing for P (i.e., $\Delta \vdash P$ holds), there is an output μ with subject a s.t. $P \xRightarrow{\mu} P'$, and such a transition is observable under the typing Δ . The meaning of 'observable under a typing' will depend on the specific type system adopted; in the case of the plain sorting, all transitions are observable.

Having typed processes, in the definition of barbed equivalence we may only test processes with contexts that respect the typing of the processes.

Definition 1. C is a Γ/Δ context if $\Gamma \vdash C$ holds, using the typing for the processes plus the rule $\frac{}{\Delta \vdash [\cdot]}$ for the hole.

Similarly, P is a Δ -process if $\Delta \vdash P$. We assume (as in usual Subject-Reduction properties for type systems) that typing is invariant under reduction. This is the case for the type systems we present in this paper.

Definition 2 (Barbed bisimulation, equivalence, and congruence). *Barbed Δ -bisimulation* is the largest symmetric relation $\approx^\Delta$ on Δ -processes s.t. $P \approx^\Delta Q$ implies:

1. whenever $P \longrightarrow P'$ then there exists Q' such that $Q \Rightarrow Q'$ and $P' \approx^\Delta Q'$;
2. for each name a , $\Delta \models P \Downarrow_a$ iff $\Delta \models Q \Downarrow_a$.

Two Δ -processes P and Q are *barbed equivalent at Δ* , written $P \simeq^\Delta Q$, if for each Γ/Δ static context E it holds that $E[P] \approx^\Gamma E[Q]$. *Barbed congruence at Δ* , $\cong^\Delta$, is defined in the same way but employing all Γ/Δ contexts (rather than only the static ones).

Barbed equivalence in the plain (untyped) $\lambda\pi$, $\simeq$, can be proved to coincide with the ordinary labeled early asynchronous bisimilarity, on image-finite processes, exploiting the n -approximants of the labeled equivalences. We recall that the class of *image-finite processes* is the largest subset $\mathcal{I}$ of processes that is derivation closed and s.t. $P \in \mathcal{I}$ implies that, for all μ , the set $\{P' \mid P \xRightarrow{\mu} P'\}$, quotiented by alpha conversion, is finite. Similarly, we present below the *early asynchronous bisimulation*, named so due to the second clause for input transitions; we simply call it *bisimulation*, and do so for the coinductive equivalences we consider in the paper, which are all asynchronous and use the early labeled transitions.

Definition 3 (Bisimulation). A relation $\mathcal{R}$ on processes is a *bisimulation* if whenever $P \mathcal{R} Q$ and $P \xrightarrow{\mu} P'$, then one of these two clauses holds:

1. there is Q' such that $Q \xRightarrow{\mu} Q'$ and $P' \mathcal{R} Q'$;
2. $\mu = a(\tilde{b})$ and there is Q' such that $Q|\bar{a}(\tilde{b}) \Rightarrow Q'$ and $P' \mathcal{R} Q'$.

Moreover the converse holds too, on the transitions from Q . *Bisimilarity*, $\approx$, is the largest bisimulation.

Theorem 1 ([18]). *On image-finite processes, relations $\simeq$ and $\approx$ coincide.*

$$\begin{array}{c}
\text{I-ACT} \frac{1 \vdash_s P}{1 \vdash_s u(\tilde{a}).P} \quad \text{I-INA} \frac{1 \vdash_s P}{0 \vdash_s x(\tilde{a}).P} \quad \text{I-INA!} \frac{1 \vdash_s P}{0 \vdash_s !x(\tilde{a}).P} \quad \text{O-ACT} \frac{}{1 \vdash_s \bar{x}(\tilde{a})} \quad \text{O-INA} \frac{}{0 \vdash_s \bar{u}(\tilde{a})} \\
\\
\text{RES} \frac{\eta \vdash_s P}{\eta \vdash_s (\nu a)P} \quad \text{NIL} \frac{}{0 \vdash_s \mathbf{0}} \quad \text{PAR} \frac{\eta_1 \vdash_s P \quad \eta_2 \vdash_s Q}{\eta_1 + \eta_2 \vdash_s P \mid Q} \quad \eta_1 + \eta_2 \leq 1 \quad \text{SUM} \frac{\eta \vdash_s G_1 \quad \eta \vdash_s G_2}{\eta \vdash_s G_1 + G_2} \\
\\
\text{TAU} \frac{\eta \vdash_s P}{\eta \vdash_s \tau.P} \quad \text{MAT} \frac{0 \vdash_s G}{0 \vdash_s [a = b]G}
\end{array}$$

Fig. 3. The typing rules for sequentiality.

3. Sequentiality

In this section we study sequentiality. We first formalize it by means of a type system, and then we examine its impact on behavioral equivalence.

3.1. Type system

As mentioned in Section 1, intuitively, sequentiality ensures us that at any time at most one interaction can occur in a system; i.e., there is a single computation thread. A process that holds the thread decides what the next interaction can be. It does so by offering a single particle (input or output) that controls the thread. The process may offer multiple particles, but only one of them may control the thread. The control on the thread attached to a particle is determined by the subject name of that particle. A given name may exercise the control on the thread either in output or in input; in the former case we say that the name is *output-controlled*, in the latter case the name is *input-controlled*. For instance, suppose that x, y, z are output-controlled and u, v are input-controlled. Then the following process correctly manages the thread and will indeed be typable in our type system:

$$P \stackrel{\text{def}}{=} u.(\bar{x}|y.\bar{x})|z.\bar{y}|\bar{v}$$

The initial particles in P are $u, z, \bar{v}$; however only u controls the thread, as z is output-controlled and v is input-controlled. When the input at u is consumed, the new particles $\bar{x}, y$ are available, where $\bar{x}$ now controls the thread, as both names x, y are output-controlled. An external process that consumes the particle $\bar{x}$ will acquire the control over the thread. For instance, a process such as $Q \stackrel{\text{def}}{=} \bar{u}|x.Q'$ initially does not hold the thread; in the parallel composition $P|Q$, after the two interactions at u and x , the control on the thread will be acquired by Q' :

$$P|Q \longrightarrow (y.\bar{x}|z.\bar{y}|\bar{v})|Q'$$

Now Q' will decide on the next interaction; for instance, it may offer an output at y or z , or an input at v . It may only offer one of these, though it may offer other particles that do not control the thread.

Notation In the remainder, x, y, z range over output-controlled names, u, v, w over input-controlled names; we recall that a, b, c range over the set of all names.

The type system for sequentiality is presented in Fig. 3. Judgements are of the form $\eta \vdash_s P$, for $\eta \in \{0, 1\}$. A judgement $1 \vdash_s P$ indicates that P owns the thread, i.e., P is *active*, and $0 \vdash_s P$ otherwise, i.e., P is *inactive*. We recall that we only present the additional typing constraints given by sequentiality, assuming the existence of a sorting under which all processes are well-typed (thus the fully-fledged typing judgements would be the form $\Delta; \eta \vdash_s P$, rather than $\eta \vdash_s P$).

Some remarks on the rules in Fig. 3: the continuation of an input always owns control on the thread; the input itself may or may not have the control (rules I-ACT and I-INA). A τ -prefix is neutral w.r.t. the thread. The rule for parallel composition makes sure that the control on the thread is granted to only one of the components; in contrast, in the rule for sum, the control is maintained for both summands. Operators $\mathbf{0}$ and match cannot own the thread; this makes sure that the thread control is always exercised.

We present some behavioral properties that highlight the meaning of sequentiality. A reduction $P \xrightarrow{\tau} P'$ is an *interaction* if it has been obtained from a communication between an input and an output (formally, its derivation in the LTS of Fig. 2 uses rule ACOMM). A pair of an unguarded input and an unguarded output at the same name form an *interaction redex*. In a sequential system, one may not find two disjoint interaction redexes. The propositions below are proved by a straightforward inspection on the typing rules.

Proposition 1. *Whenever $\eta \vdash_s P$, there exists no $P_1, P_2, \tilde{a}$ such that $P \equiv (\nu \tilde{a})(P_1|P_2)$ with $P_1 \xrightarrow{\tau} P'_1$ and $P_2 \xrightarrow{\tau} P'_2$, and both these transitions are interactions.*

An inactive process may not perform interactions.

Proposition 2. *If $0 \vdash_s P$, then P does not have an interaction redex; i.e., there is no P' with $P \xrightarrow{\tau} P'$ and this transition is an interaction.*

An inactive process may however perform τ -reductions, notably to resolve internal choices. In other words such internal choices represent internal matters for a process, orthogonal with respect to the overall interaction thread. The possibility for inactive processes to accommodate internal choices will be important in our completeness proof. However, an inactive process may only perform a finite number of τ -reductions, intuitively because replications in the language are input-guarded. A process P is *divergent* if it can perform an infinite sequence of reductions, i.e., there are $P_1, P_2, \dots$, with $P \longrightarrow P_1 \longrightarrow P_2 \dots P_n \longrightarrow \dots$.

Proposition 3. *If $0 \vdash_s P$ then P is not divergent.*

In contrast, an active process may be divergent, through sequences of reductions containing infinitely many interactions.

Sequentiality imposes constraints on the interactions that a ‘legal’ (i.e., well-typed) context may undertake with a process. For the definition of barbed bisimulation and equivalence we must therefore define the meaning of observability. The following definition of *type-allowed transitions* shows what such ‘legal’ interactions can be.

Definition 4 (Type-allowed transitions). We write $\eta \models_s P \xrightarrow{\mu} P'$ if $\eta \vdash_s P$, and $P \xrightarrow{\mu} P'$, and one of the following clauses holds:

1. $\eta = 0$
2. $\mu = \tau$
3. $\eta = 1$ and $\mu = u\langle\tilde{a}\rangle$ for some $u, \tilde{a}$ or $\mu = (\nu\tilde{a})\bar{x}\langle\tilde{b}\rangle$ for some $\tilde{a}, x, \tilde{b}$.

Clause (1) says that all interactions between an inactive process and the context are possible; this holds because the context is active and may therefore decide on the next interaction with the process. Clause (2) says that internal reductions may always be performed. Clause (3) says that the only visible actions observable in active processes are those carrying the thread; this holds because the observer is inactive, and it is therefore up to the process to decide on the next interaction.

We now examine how typing evolves under legal actions. We recall that x stands for an output-controlled name.

Definition 5. We write $[\eta; P] \xrightarrow{\mu} [\eta'; P']$ when $\eta \models_s P \xrightarrow{\mu} P'$ and:

1. if $\mu = x\langle\tilde{a}\rangle$, then $\eta' = 1$.
2. if $\mu = (\nu\tilde{a})\bar{x}\langle\tilde{b}\rangle$, then $\eta' = 0$.
3. otherwise $\eta' = \eta$.

Lemma 1. *If $\eta \vdash_s P$ and $P \xrightarrow{\mu} P'$, with either $\mu = u\langle\tilde{a}\rangle$ for some $u, \tilde{a}$ or $\mu = (\nu\tilde{a})\bar{x}\langle\tilde{b}\rangle$ for some $\tilde{a}, x, \tilde{b}$, then $\eta = 1$.*

Proof. Recall that $u\langle\tilde{a}\rangle.P$ and $\bar{x}\langle\tilde{b}\rangle$ are the prefixes typed with $\eta = 1$ and must be the base case for the derivation of the transition $P \xrightarrow{\mu} P'$. The statement follows by a simple induction on the transition, we show that η never decreases to 0. $\square$

This lemma essentially means that the mention of $\eta = 1$ in Clause (3) of Definition 4 is not required. We can now prove subject reduction.

Theorem 2 (Subject Reduction).

If $\eta \vdash_s P$ and $[\eta; P] \xrightarrow{\mu} [\eta'; P']$ then $\eta' \vdash_s P'$.

Proof. We go by induction on the transition $P \xrightarrow{\mu} P'$. We discuss the main cases; the remaining ones are simpler.

- Case **APAR** : The assumption can be rewritten as $[\eta; P_1 | P_2] \xrightarrow{\mu} [\eta'; P'_1 | P'_2]$, with $P_1 \xrightarrow{\mu} P'_1$. By rule **PAR**, there exist η_1, η_2 such that $\eta_i \vdash_s P_i$ for $i \in \{1, 2\}$ and $\eta = \eta_1 + \eta_2$.

The first step is to note that $\eta_1 \models_s P_1 \xrightarrow{\mu} P'_1$, which is easily proved for all three cases of Definition 4, with the modified version of Clause (3) via Lemma 1. Thus, there always exists a η'_1 such that $[\eta_1; P_1] \xrightarrow{\mu} [\eta'_1; P'_1]$, where η'_1 depends on the corresponding case for μ in Definition 5. By induction, it implies $\eta'_1 \vdash_s P'_1$. By noticing that we must have $0 \vdash_s P_2$ in the first two cases of Definition 5, we can prove that $\eta' = \eta'_1 + \eta_2$ and thus $\eta' \vdash_s P'_1 | P'_2$.

- Case **ACOMM** : The assumption can be rewritten as $[\eta; P_1 | P_2] \xrightarrow{\tau} [\eta'; (\nu\tilde{c})(P'_1 | P'_2)]$, with $P_1 \xrightarrow{a\langle\tilde{b}\rangle} P'_1$ and $P_2 \xrightarrow{(\nu\tilde{c})\bar{a}\langle\tilde{b}\rangle} P'_2$. If a is an output-controlled name x , then by Lemma 1, $1 \vdash_s P_2$ and thus by rule **PAR**, $0 \vdash_s P_1$. On the other hand, if a is an input-controlled name u , then $1 \vdash_s P_1$ and $0 \vdash_s P_2$. In both cases, we have $\eta_1 \vdash_s P_1$, $\eta_2 \vdash_s P_2$ with both $\eta_1 \models_s P_1 \xrightarrow{a\langle\tilde{b}\rangle} P'_1$ and $\eta_2 \models_s P_2 \xrightarrow{(\nu\tilde{c})\bar{a}\langle\tilde{b}\rangle} P'_2$.

Thus there always exists a η'_1 such that $[\eta_1; P_1] \xrightarrow{a\langle\tilde{b}\rangle} [\eta'_1; P'_1]$, and thus by induction, it implies $\eta'_1 \vdash_s P'_1$. A simple case distinction shows that $\eta'_1 = 1$. Similarly, we can prove that $[\eta_2; P_2] \xrightarrow{(\nu\tilde{c})\bar{a}\langle\tilde{b}\rangle} [0; P'_2]$, and thus $0 \vdash_s P'_2$ by induction. Finally, by rule **PAR** and **RES**, we have $1 \vdash_s (\nu\tilde{c})(P'_1 | P'_2)$.

$\square$

Weak type-allowed transitions are defined as expected, exploiting the invariance of typing under reductions: $\eta \models_s P \xRightarrow{\mu} P'$ holds if there are P_0, P_1 with $P \xRightarrow{\mu} P_0$, $\eta \models_s P_0 \xrightarrow{\mu} P_1$ and $P_1 \xRightarrow{\mu} P'$.

3.2. Behavioral equivalence

3.2.1. Sequential bisimilarity

To tune [Definition 2](#) of barbed bisimulation and equivalence to the setting of sequentiality, we have to specify the meaning of observables. An observable $\eta \models_s P \Downarrow_a$ holds if there are P' and an output action μ such that $\eta \models_s P \xrightarrow{\mu} P'$ and the subject of μ is a . Following [Definition 1](#), in barbed equivalence, the legal contexts are the η/η' static contexts. To remind ourselves of the sequentiality constraint, we write barbed equivalence at η as $\simeq^\eta$, and sometimes call it *sequential barbed equivalence at η* . Thus $P \simeq^\eta Q$ holds if $\eta \vdash_s P, Q$ and $E[P] \approx^{\eta'} E[Q]$, for any η' and any η'/η static context E . We are now ready to define the labeled bisimilarity to be used on sequential processes, which is our main proof technique for barbed equivalence. A *typed process relation* is a set of triplets (η, P, Q) with $\eta \vdash_s P, Q$.

Definition 6 (Sequential Bisimulation). A typed process relation $\mathcal{R}$ is a *sequential bisimulation* if whenever $(\eta, P, Q) \in \mathcal{R}$ and $[\eta; P] \xrightarrow{\mu} [\eta'; P']$, then one of the two following clauses holds:

1. there is Q' such that $Q \xrightarrow{\hat{\mu}} Q'$ and $(\eta', P', Q') \in \mathcal{R}$;
2. $\mu = a(\tilde{b})$ and there is Q' such that $Q|\tilde{a}(\tilde{b}) \Rightarrow Q'$ with $(\eta', P', Q') \in \mathcal{R}$.

Moreover, the converse of (1) and (2) holds on the transitions from Q . Processes P and Q are *sequentially bisimilar at η* , written $P \approx_s^\eta Q$, if $(\eta, P, Q) \in \mathcal{R}$ for some sequential bisimulation $\mathcal{R}$.

In clause (2), $Q|\tilde{a}(\tilde{b})$ must be well-typed, whether the name a is input- or output-controlled. Clauses (1) and (2) are the same as for ordinary bisimilarity $\approx$ ([Definition 3](#)); typing however prevents certain transitions to be considered as challenge transitions in the bisimulation game. Thus the resulting bisimilarity becomes coarser.

Ordinary bisimilarity, restricted to processes having the same typing, is included in the sequential one. So it is also a proof technique for barbed equivalence at η . Indeed, from any bisimulation, we may construct a sequential bisimulation by taking the pairs of well-typed processes in relation.

Proposition 4. For $\eta \vdash_s P, Q$, if $P \approx Q$ then also $P \approx_s^\eta Q$.

The inclusion is strict: all the examples in [Section 3.3](#) fail to hold for $\approx$.

The following theorem states the soundness of the labeled bisimilarity $\approx_s^\eta$. As usual, its proof relies on the preservation of $\approx_s^\eta$ under parallel composition, which requires some care in order to enforce sequentiality, and which is ensured by typability.

Theorem 3 (Soundness). If $P \approx_s^\eta Q$, then $P \simeq^\eta Q$.

Proof. Since sequential bisimilarity is included in sequential barbed bisimilarity, the proof amounts to showing that sequential bisimilarity is closed by typed static contexts, i.e., that the obtained closure, noted $\mathcal{R}$, is a sequential bisimulation. Take P, Q, R such that $P \approx_s^\eta Q$, $\eta' \vdash_s R$ and $\eta + \eta' \leq 1$. We must show that the triple $(\eta + \eta', (\nu\tilde{c})(P|R), (\nu\tilde{c})(Q|R)) \in \mathcal{R}$ satisfies the bisimulation property. Since restrictions do not interact with the type system, we focus on a triple without restriction. Take a typed transition $[\eta + \eta'; P|R] \xrightarrow{\mu} [\eta''; P']$, we have three cases:

- $[\eta + \eta'; P|R] \xrightarrow{\mu} [\eta''; P_2|R]$ (rule **PAR**): With the same reasoning as in the proof of [Theorem 2](#), we have that $[\eta; P] \xrightarrow{\mu} [\eta_1; P_2]$ with $\eta_1 + \eta' = \eta''$. As $P \approx_s^\eta Q$, we have:
 - Either there exists Q_2 such that $Q \xrightarrow{\hat{\mu}} Q_2$ and $P_2 \approx_s^{\eta_1} Q_2$. This means we have $\eta_1 \vdash_s Q_2$ and $Q|R \xrightarrow{\hat{\mu}} Q_2|R$ and $\eta_1 + \eta' \vdash_s Q_2|R$, so $(\eta'', P_2|R, Q_2|R) \in \mathcal{R}$.
 - Or $\mu = a(\tilde{b})$ and there exists Q_2 such that $Q|\tilde{a}(\tilde{b}) \Rightarrow Q_2$ and $P_2 \approx_s^{\eta_1} Q_2$. This means we have $\eta_1 \vdash_s Q_2$ and $\eta_1 \vdash_s Q|\tilde{a}(\tilde{b})$ by subject reduction. Thus, we have $Q|R|\tilde{a}(\tilde{b}) \Rightarrow Q_2|R$ and $\eta_1 + \eta' \vdash_s Q_2|R$ so $(\eta'', P_2|R, Q_2|R) \in \mathcal{R}$.
- $[\eta + \eta'; P|R] \xrightarrow{\mu} [\eta''; P|R_2]$ (rule **PAR**): Similarly to the previous case, we have $[\eta; R] \xrightarrow{\mu} [\eta'_1; R_2]$ with $\eta + \eta'_1 = \eta''$. Thus, by subject reduction, we have $\eta'_1 \vdash_s R_2$ and thus $\eta'' \vdash_s Q|R_2$, so $(\eta'', P|R_2, Q|R_2) \in \mathcal{R}$. We can then conclude as $Q|R \xrightarrow{\mu} Q|R_2$.
- $[\eta + \eta'; P|R] \xrightarrow{\tau} [\eta''; (\nu\tilde{c})(P_2|R_2)]$ (rule **COMM**): We show the case where the output comes from P , i.e., $P \xrightarrow{(\nu\tilde{c})\tilde{a}(\tilde{b})} P_2$ and $R \xrightarrow{a(\tilde{b})} R_2$. With the same reasoning as in the proof of [Theorem 2](#), we have $[\eta; P] \xrightarrow{(\nu\tilde{c})\tilde{a}(\tilde{b})} [\eta_2; P_2]$, and $[\eta'; R] \xrightarrow{a(\tilde{b})} [\eta'_2; R_2]$ with $\eta_2 + \eta'_2 = \eta''$. As $P \approx_s^\eta Q$, there exists Q_2 such that $Q \xrightarrow{(\nu\tilde{c})\tilde{a}(\tilde{b})} Q_2$ and $P_2 \approx_s^{\eta_2} Q_2$. This means we have $\eta_2 \vdash_s Q_2$ and $Q|R \xrightarrow{\tau} (\nu\tilde{c})(Q_2|R_2)$ and $\eta_2 + \eta'_2 \vdash_s (\nu\tilde{c})(Q_2|R_2)$, so $(\eta'', (\nu\tilde{c})(P_2|R_2), (\nu\tilde{c})(Q_2|R_2)) \in \mathcal{R}$.

□

[Theorem 3](#) allows us to use the labeled bisimilarity $\approx_s^\eta$ as a proof technique for typed barbed equivalence. This proof technique is also complete, assuming only output-controlled names (i.e., the thread may only be exercised by output particles, not by the input ones), as we show below.

3.2.2. Completeness for output-controlled names

For the proof of completeness, we introduce the notion of *singular* process. Intuitively, a singular process always keeps the thread. Formally, the set of singular processes is the largest set $\mathcal{T}$ of processes such that for all $P \in \mathcal{T}$, we have $1 \vdash_s P$ and whenever $1 \models_s P \xrightarrow{\mu} P'$, then $P' \in \mathcal{T}$. A simple example of a singular process is $0_1 \stackrel{\text{def}}{=} (\nu z)\bar{z}$, which is an active process without transitions. Singular processes are simply described as the equivalence class of 0_1 .

Lemma 2. *P is singular iff $P \approx_s^1 0_1$.*

Because a singular process always keeps the thread, the presence of other components is irrelevant, and cannot be observed.

Lemma 3. *For any static $1/1$ context E and singular process P , we have $E[P] \approx_s^1 P$.*

For processes with only output-controlled names, we may use a coarser equivalence to tell whether a process is singular.

Lemma 4. *If P only contains output-controlled names, then P is singular iff $P \approx_s^1 0_1$.*

This result does not extend to general processes. For instance, we have $u.\bar{x} \approx_s^1 0_1$ but $u.\bar{x} \not\approx_s^1 0_1$. We can now describe the completeness proof. While the overall structure of the proof is standard, the technical details are specific to sequentiality. As usual, we rely on a stratification of bisimilarity and approximants.

Definition 7 (Approximants of sequential bisimilarity). We define a sequence $(\approx_s^{\eta,n})_{n \geq 0}$ of typed relations:

1. $P \approx_s^{\eta,0} Q$ if $\eta \vdash_s P, Q$.
2. For $1 \leq i$, relation $\approx_s^{\eta,i}$ is defined by: $P \approx_s^{\eta,i} Q$ if whenever $[\eta; P] \xrightarrow{\mu} [\eta'; P']$, one of these two clauses holds:
 - there is Q' such that $Q \xRightarrow{\hat{\mu}} Q'$ and $P' \approx_s^{\eta',i-1} Q'$;
 - $\mu = a(b)$ and there is Q' such that $Q|a(b) \Rightarrow Q'$ and $P' \approx_s^{\eta',i-1} Q'$,
 and symmetrically for the transitions of Q .
3. Then $P \approx_s^{\eta,\omega} Q$ if $P \approx_s^{\eta,i} Q$ for all i .

Notice that $\approx_s^{\eta,0} \supseteq \approx_s^{\eta,1} \supseteq \dots \supseteq \approx_s^{\eta,n} \supseteq \dots \supseteq \approx_s^{\eta,\omega} \supseteq \approx_s^{\eta}$.

On image-finite processes, sequential bisimilarity, $\approx_s^{\eta}$, coincides with the intersection of its approximants.

Lemma 5. *If P, Q are image-finite, $P \approx_s^{\eta,\omega} Q$ iff $P \approx_s^{\eta} Q$.*

The main ingredient for proving completeness is given by the following proposition.

Proposition 5. *compleseq Given $n \geq 0$, suppose $\eta \vdash_s P, Q$, and $P \not\approx_s^{\eta,n} Q$. Suppose further that both P and Q only use output-controlled names and are image-finite. Then there exists R such that for any fresh name z , any $\tilde{x}, S$ with $\tilde{x} \subseteq \text{fn}(P) \cup \text{fn}(Q) \subseteq S$, if we define*

$$\bar{R} \stackrel{\text{def}}{=} \begin{cases} \bar{z}|z.R & \text{when } 0 \vdash_s P, Q \\ R + \sum_{y \in S} y(\tilde{y}).\bar{z} & \text{when } 1 \vdash_s P, Q \end{cases}$$

and $E_0 \stackrel{\text{def}}{=} (\nu \tilde{x})([\cdot]| \bar{R})$,
we have the following:

1. either $E_0[P'] \not\approx_s^1 E_0[Q]$ for all P' such that $P \Rightarrow P'$
2. or $E_0[P] \not\approx_s^1 E_0[Q']$ for all Q' such that $Q \Rightarrow Q'$.

When $1 \vdash_s P, Q$, we define R as a guarded process, in order for the sum in the definition of $\bar{R}$ to make sense. The full proof is given in [Appendix A](#). We sketch here the main aspects, specific to the case of sequential bisimilarity. We reason by contradiction to establish that $E_0[P]$ and $E_0[Q]$ are not barbed bisimilar. To respect typing, the definition of E_0 requires some care.

The case when $\eta = 0$ (tested processes are inactive) is rather standard: the context E_0 is of the form $(\nu \tilde{x})([\cdot]| \bar{z}|z.R)$, for some fresh z , and some “tester process” R . The barb at z allows us to detect when the tested process interacts with R .

The delicate case is when $\eta = 1$ (tested processes are active): the context must be inactive and hence cannot have an unguarded output at z . We use in this case a $1/1$ context of the form $E_0 \stackrel{\text{def}}{=} (\nu \tilde{x})([\cdot]| G_R + G)$. Process G_R is the tester process, and G is $\sum_{y \in S} y(\tilde{y}).\bar{z}$, defined for some fresh z and some set S containing $\text{fn}(P) \cup \text{fn}(Q)$. G satisfies the following property: for any P_0 and for any x , if $1 \models_s P_0 \Downarrow_{\bar{x}}$ then $1 \models_s P_0|G \Downarrow_{\bar{z}}$. Thus, as soon as P_0 exhibits some barb, we have $1 \models_s E_0[P_0] \Downarrow_{\bar{z}}$, and P_0 cannot interact with R without removing the barb at z , which allows us to reason as in the case $\eta = 0$.

The proof schema above relies on having P_0 active and exhibiting a barb, in order to be able to detect the interactions. This condition is not fulfilled for singular processes, which must be handled separately. Since we restrict to output-controlled names, we can use [Lemma 4](#). If on the contrary input-controlled names are allowed, P_0 may fail to exhibit any barb while not being singular, which prevents us from using the same reasoning.

Theorem 4 (Completeness on output-controlled names). *For all image-finite processes P, Q that only use output-controlled names, and for all η , if $P \simeq^\eta Q$ then $P \approx_s^\eta Q$.*

Proof. We suppose $P \simeq^\eta Q$ and P, Q are image-finite. We reason by contradiction: if $P \not\approx_s^\eta Q$, then by Lemma 5, $P \not\approx_s^{\eta, \omega} Q$. Thus there exists n such that $P \not\approx_s^{\eta, n} Q$. By Proposition 5, there exists some static context E such that $E[P] \not\approx_s^{\eta, n} E[Q]$. Hence $P \not\approx_s^\eta Q$, which is absurd. $\square$

3.3. Examples

With respect to ordinary bisimilarity, in sequential bisimilarity ($\approx_s^\eta$) fewer challenges are allowed. This may both mean that certain processes, otherwise distinguishable, become equal, and that certain equalities are simpler to prove because the state space of the processes to be examined is reduced. We present some equalities of the first kind (valid for $\approx_s^\eta$ only). In Section 3.3.2, we also show a refinement of $\approx_s^\eta$ which is useful for reasoning about references.

3.3.1. Basic examples

Example 1. An ungarded occurrence of an input at an input-controlled name becomes the only observation that can be made in a process. This yields the following equalities

$$\begin{array}{lll} u.P|x.Q & \approx_s^1 & u.(P|x.Q) \\ u.P|\bar{v} & \approx_s^1 & u.(P|\bar{v}) \end{array} \quad \text{for } u \neq v$$

For the first equality, one shows that

$$\{(1, u.P|x.Q, u.(P|x.Q))\} \cup I,$$

where I is the (typed) identity relation, is a sequential bisimulation. One proves the second equality similarly.

By sequentiality, it is also not possible to access parts of processes that would require a simultaneous/parallel reception. The following example illustrates this.

Example 2. Consider the process

$$P \stackrel{def}{=} (\nu y', z')(!x.(z'.\bar{z}|y')|!y.\bar{z}).$$

The output at z becomes observable if both an input at x and an input at y are consumed, so that the internal reduction at z' can take place. However the input at x acquires the thread, preventing a further immediate input at y ; similarly for the input at y . Indeed we have $P \approx_s^0 x.0_1 + y.0_1 \stackrel{def}{=} Q$ (we recall that $0_1 = (\nu z)\bar{z}$). To prove the equality, we can use $\{P, Q\} \cup \approx_s^1$, which is easily proved to be a sequential bisimulation. The derivatives of P and Q are singular (and stable, i.e., unable to reduce further) processes; therefore they are in $\approx_s^1$, as discussed in Lemma 4.

Under the ordinary bisimilarity, P and Q are distinguished because the sequence of transitions $P \xrightarrow{x} y \xrightarrow{\tau} \bar{z}$ cannot be matched by Q .

Example 3. This example informally discusses why sequentiality can help reducing the number of pairs of processes to examine in a bisimulation proof. Suppose we wish to prove the equality between the two processes

$$\begin{array}{ll} P_1 & \stackrel{def}{=} \bar{v}|y.R|u.Q \\ P_2 & \stackrel{def}{=} \bar{v}|u.Q|y.R \end{array}$$

and such processes are typable under the sequentiality system. The difference between P_1 and P_2 comes from commutativity of parallel composition. We may therefore use ordinary bisimilarity, as this is a sound proof technique for typed barbed equivalence. One may prove, more generally, that parallel composition is commutative and derive the equality. However suppose we wish to use the bisimulation method, concretely, on P and Q . The two processes have 3 initial transitions (with labels $\bar{v}$, y , and u), and the subtrees of transitions emanating from the such derivatives have similar size. Under ordinary bisimulation, we have to examine all the states in the 3 subtrees. By taking $R = \bar{x}$ and $Q = \bar{z}$, it would lead to the following relation with 18 pairs:

$$\begin{aligned} & \{(\bar{v}|y.\bar{x}|u.\bar{z}, \bar{v}|u.\bar{z}|y.\bar{x}), (0|y.\bar{x}|u.\bar{z}, 0|u.\bar{z}|y.\bar{x}), (\bar{v}|\bar{x}|u.\bar{z}, \bar{v}|u.\bar{z}|\bar{x}), (\bar{v}|0|u.\bar{z}, \bar{v}|u.\bar{z}|0), \\ & (\bar{v}|y.\bar{x}|\bar{z}, \bar{v}|\bar{z}|y.\bar{x}), (\bar{v}|y.\bar{x}|0, \bar{v}|0|y.\bar{x}), (0|\bar{x}|u.\bar{z}, 0|u.\bar{z}|\bar{x}), (0|0|u.\bar{z}, 0|u.\bar{z}|0), \\ & (0|y.\bar{x}|\bar{z}, 0|\bar{z}|y.\bar{x}), (0|y.\bar{x}|0, 0|0|y.\bar{x}), (\bar{v}|\bar{x}|\bar{z}, \bar{v}|\bar{z}|\bar{x}), (\bar{v}|0|\bar{z}, \bar{v}|\bar{z}|0), (\bar{v}|\bar{x}|0, \bar{v}|0|\bar{x}), \\ & (\bar{v}|0|0, \bar{v}|0|0), (0|\bar{x}|\bar{z}, 0|\bar{z}|\bar{x}), (0|0|\bar{z}, 0|\bar{z}|0), (0|\bar{x}|0, 0|0|\bar{x}), (0|0|0, 0|0|0)\} \end{aligned}$$

Under sequential bisimulation, however, only the input at u is observable (it is the only one carrying the thread), thus removing 2 of the 3 initial subtrees. In our minimal example, it would reduce it to only 11 triples:

$$\begin{aligned} & \{(1, \bar{v}|y.\bar{x}|u.\bar{z}, \bar{v}|u.\bar{z}|y.\bar{x}), (1, \bar{v}|y.\bar{x}|\bar{z}, \bar{v}|\bar{z}|y.\bar{x}), (0, \bar{v}|y.\bar{x}|0, \bar{v}|0|y.\bar{x}), \\ & (0, 0|y.\bar{x}|\bar{z}, 0|0|y.\bar{x}), (1, 0|y.\bar{x}|\bar{z}, 0|\bar{z}|y.\bar{x}), (1, \bar{v}|\bar{x}|0, \bar{v}|0|\bar{x}), (1, \bar{v}|0|\bar{z}, \bar{v}|\bar{z}|0), \\ & (0, \bar{v}|0|0, \bar{v}|0|0), (1, 0|\bar{x}|0, 0|0|\bar{x}), (1, 0|0|\bar{z}, 0|\bar{z}|0), (0, 0|0|0, 0|0|0)\} \end{aligned}$$

Further pruning may be possible later on, again exploiting the fact that under sequentiality only certain transitions are observable.

$$\begin{array}{c}
\frac{}{\ell \vdash_r \bar{\ell}\langle n \rangle} \quad \frac{\ell \vdash_r P}{\emptyset \vdash_r \ell(m).P} \quad \frac{S, \ell \vdash_r P}{S \vdash_r (\nu \ell)P} \quad \frac{S_1 \vdash_r P \quad S_2 \vdash_r Q}{S_1 \uplus S_2 \vdash_r P|Q} \quad \frac{S \vdash_r P}{S \vdash_r (\nu a)P} \\
\frac{}{\ell \vdash_r \bar{\ell}\langle n \rangle} \quad \frac{\ell \vdash_r P}{\emptyset \vdash_r \ell(m).P} \quad \frac{S, \ell \vdash_r P}{S \vdash_r (\nu \ell)P} \quad \frac{S_1 \vdash_r P \quad S_2 \vdash_r Q}{S_1 \uplus S_2 \vdash_r P \mid Q} \quad \frac{S \vdash_r P}{S \vdash_r (\nu a)P}
\end{array}$$

Fig. 4. Rules for references.

3.3.2. Examples with references

We now consider a few examples involving references. For this, we use the standard encoding of references into $\mathcal{A}\pi$, and we enhance the bisimilarity for sequentiality so to take references into account.

We use $n, m, \dots$ to range over the entities stored in references (which can be names or values belonging to a first-order data type like booleans and integers) and placeholders for them. Name ℓ is used to represent a reference.

In $\mathcal{A}\pi$, a reference ℓ holding a value n is represented as an output particle $\bar{\ell}\langle n \rangle$. A process that contains a reference ℓ should have, at any time, exactly one unguarded output at ℓ , meaning that at any time the reference has a unique value. We say that in this case ℓ is *accessible*. The read and write operations on ℓ are written as follows:

$$\begin{aligned}
\text{re}_\ell(m).R &\stackrel{\text{def}}{=} \ell(m).(\bar{\ell}\langle m \rangle | R) \\
\text{wr}_\ell\langle n \rangle.R &\stackrel{\text{def}}{=} \ell(m').(\bar{\ell}\langle n \rangle | R) \text{ for } m' \notin \text{fn}(R)
\end{aligned}$$

Thus a name ℓ used to encode a reference is input-controlled, as an action on a reference is represented by an input at ℓ | we use ℓ rather than $u, v, \dots$ to stress the fact that names used to represent references obey constraints that go beyond input-control.

Proof techniques for the representation of references in $\mathcal{A}\pi$ have been studied in [17]. Adopting them requires enhancing our type system with information about references, which simply consists in declaring which names represent references. In Fig. 4, we give the main rules for defining $S \vdash_r P$ where S is the set of accessible reference names in P . The other rules are standard having no accessible references names neither in premise nor in conclusion. In the definition of barbed equivalence, the main constraint is that the tested context should make sure that all existing reference names are accessible. To reason about references, several definitions of labeled bisimilarity are presented in [17], varying on the forms of constraints imposed on transitions. Here we only import the simplest such constraint: we forbid observations of input transitions $P \xrightarrow{\ell(n)} P'$ at a reference name ℓ when ℓ is accessible in P (i.e., an unguarded output at ℓ occurs in P). Such a constraint represents the fact that an observer may not pretend to own a reference when the reference is accessible in the process.

Formally, with the addition of references, judgements in the type system become of the form $S; \eta \vdash_s P$, where S is a finite set of reference names, meaning that $\eta \vdash_s P$ holds and that $S \vdash_r P$ holds. The definition of type-allowed transitions, $S; \eta \models_s P \xrightarrow{\mu} P'$, is the same as before (Definition 4) with the addition, in clause (3), of the constraint

$$\text{if } \mu \text{ is an input } \ell\langle n \rangle \text{ at a reference name } \ell \text{ then } \ell \notin S. \quad (*)$$

We then define transitions of the form $[S; \eta; P] \xrightarrow{\mu} [S'; \eta'; P']$, along the lines of Definition 5. The S component evolves along interactions at reference names: ℓ is added to S along an input at ℓ or if $\ell \in \text{bn}(\mu)$ (since ℓ is now visible by the environment), while ℓ is removed from S along an output at ℓ . Finally, the definition of *sequential bisimilarity with references at* (S, η) , written $\approx_s^{S, \eta}$ is the same as that of sequential bisimilarity (Definition 6), just using $S; \eta \vdash_s P, Q$ and $[S; \eta; P] \xrightarrow{\mu} [S'; \eta'; P']$, in place of $\eta \vdash_s P, Q$ and $[\eta; P] \xrightarrow{\mu} [\eta'; P']$.

It is straightforward to extend the soundness proof for sequential bisimilarity w.r.t. barbed equivalence (Theorem 3) to the case of sequential bisimilarity with references. We refer to [17] for further details on proof techniques for references in $\mathcal{A}\pi$.

Example 4. This example shows that reading or writing on a global reference is not subject to interferences from the outside, as these operations require the thread:

$$\begin{aligned}
\bar{\ell}\langle n \rangle | \text{re}_\ell(m).R &\approx_s^{\ell;1} \bar{\ell}\langle n \rangle | R\{n/m\} \\
\bar{\ell}\langle n \rangle | \text{wr}_\ell\langle m \rangle.R &\approx_s^{\ell;1} \bar{\ell}\langle m \rangle | R
\end{aligned}$$

Indeed, in each law, if P (resp. Q) is the process on the left-hand (resp. right-hand) side, then the relation $\{(\ell; 1), P, Q\} \cup I$ is a sequential bisimulation, when taking the constraint $(*)$ for references into account.

Example 5 (Fetch-and-add, swap). We consider fetch-and-add and swap operations, often found in operating systems. The first, written $\text{faa}_\ell\langle n \rangle(m)$ atomically increments by n the content of the reference ℓ , and returns the original value as m ; the second, written $\text{sw}_\ell\langle n \rangle(m)$, atomically sets the content of ℓ to n and returns the original value as m :

$$\begin{aligned}
\text{faa}_\ell\langle n \rangle(m).R &\stackrel{\text{def}}{=} \ell(m).(\bar{\ell}\langle m+n \rangle | R) \\
\text{sw}_\ell\langle n \rangle(m).R &\stackrel{\text{def}}{=} \ell(m).(\bar{\ell}\langle n \rangle | R)
\end{aligned}$$

These operations may be mimicked by a combination of read and write operations (we take $m' \notin \text{fn}(R)$):

$$\begin{aligned} \text{faa}_{2\ell}\langle n \rangle(m).R &\stackrel{\text{def}}{=} \text{re}_{\ell}(m).\text{wr}_{\ell}\langle m+n \rangle.R \\ &= \ell(m).(\bar{\ell}\langle m \rangle | \ell(m').(\bar{\ell}\langle m+n \rangle | R)) \\ \text{sw}_{2\ell}\langle n \rangle(m).R &\stackrel{\text{def}}{=} \text{re}_{\ell}(m).\text{wr}_{\ell}\langle n \rangle.R \\ &= \ell(m).(\bar{\ell}\langle m \rangle | \ell(m').(\bar{\ell}\langle n \rangle | R)) \end{aligned}$$

For this mimicking to be correct, sequentiality is necessary. To see this, consider the simple case when $R \stackrel{\text{def}}{=} \bar{c}\langle m \rangle$. In the ordinary $\Lambda\pi$, processes $\text{sw}_{2\ell}\langle n \rangle(m).R$ and $\text{sw}_{2\ell}\langle n \rangle(m).R$ are distinguished, intuitively because the observer is capable of counting the two inputs and the two outputs at ℓ in $\text{sw}_{2\ell}\langle n \rangle(m).R$ (against only one in $\text{sw}_{\ell}\langle n \rangle(m).R$) and/or is capable of detecting the output $\bar{\ell}\langle m \rangle$ in $\text{sw}_{2\ell}\langle n \rangle(m).R$.

The processes are also distinguished with the proof techniques for references in [17], intuitively because, after the initial input $\ell(m)$ (whereby the processes read the content of the reference), an observer may interact with the derivative of $\text{sw}_{2\ell}\langle n \rangle(m).R$ and use its output $\bar{\ell}\langle m \rangle$ so to know the value that had been read. Such an observation is not possible with $\text{sw}_{\ell}\langle n \rangle(m).R$.

In contrast, the two processes are equal if we take sequentiality into account. That is, we have:

$$\text{sw}_{\ell}\langle n \rangle(m).R \approx_s^{\emptyset;1} \text{sw}_{2\ell}\langle n \rangle(m).R$$

This is proved by showing that the relation

$$\begin{aligned} &\cup_{m'} \{((\ell; 1), \bar{\ell}\langle n \rangle | R\{m'/m\}, \bar{\ell}\langle m' \rangle | \text{wr}_{\ell}\langle n \rangle.R\{m'/m\}) \\ &\cup I \cup \{((\emptyset; 1), \text{sw}_{\ell}\langle n \rangle(m).R, \text{sw}_{2\ell}\langle n \rangle(m).R)\} \end{aligned}$$

is a sequential bisimulation. The equivalence between $\text{faa}_{\ell}\langle n \rangle(m).R$ and $\text{faa}_{2\ell}\langle n \rangle(m).R$ is established using a similar relation.

Example 6 (Optimized access). Two consecutive read and/or write operations can be transformed into an equivalent single operation.

$$\begin{aligned} \text{wr}_{\ell}\langle n \rangle.\text{wr}_{\ell}\langle m \rangle.R &\approx_s^{\emptyset;1} \text{wr}_{\ell}\langle m \rangle.R \\ \text{wr}_{\ell}\langle n \rangle.\text{re}_{\ell}(m).R &\approx_s^{\emptyset;1} \text{wr}_{\ell}\langle n \rangle.R\{n/m\} \\ \text{re}_{\ell}(m).\text{re}_{\ell}(m').R &\approx_s^{\emptyset;1} \text{re}_{\ell}(m).R\{m/m'\} \end{aligned}$$

For the first equality, one shows that

$$\begin{aligned} &I \cup \{((\emptyset; 1), \text{wr}_{\ell}\langle n \rangle.\text{wr}_{\ell}\langle m \rangle.R, \text{wr}_{\ell}\langle m \rangle.R)\} \\ &\cup \{((\ell; 1), \bar{\ell}\langle n \rangle | \text{wr}_{\ell}\langle m \rangle.R, \bar{\ell}\langle m \rangle | R)\} \end{aligned}$$

is a sequential bisimulation. The second law is treated similarly. In both cases, the relation exhibited is finite.

For the third equality, one defines $\mathcal{R}$ as

$$\cup_n \{((\ell; 1), \bar{\ell}\langle n \rangle | \text{re}_{\ell}(m').R\{n/m\}, \bar{\ell}\langle n \rangle | R\{n/m, m'\})\}$$

Then $\{((\emptyset; 1), \text{re}_{\ell}(m).\text{re}_{\ell}(m').R, \text{re}_{\ell}(m).R\{m/m'\})\} \cup \mathcal{R} \cup I$ is a sequential bisimulation. Again, all these equalities fail under ordinary bisimilarity for the π -calculus, simply because the processes involved may perform different numbers of visible actions.

4. Well-bracketing

4.1. Type system

We now go beyond sequentiality, so to handle well-bracketing. In languages without control operators, this means that return-call interactions among terms follow a stack-based discipline. Intuitively, a well-bracketed system is a sequential system offering services. When interrogated, a service, say A , acquires the thread and is supposed to return a final result (unless the computation diverges) thus releasing the thread. During its computation, A may however interrogate another service, say B , which, upon completion of its computation, will return the result to A . In a similar manner, the service B , during its computation, may call yet another service C , and will wait for the return from C before resuming its computation. B may also delegate to C the task of returning a result to A . In any case, the ‘return’ obligation may not be thrown away or duplicated.

The implementation of this policy requires *continuation names*. For instance, when calling a service B , a process A transmits a fresh name, say p , that will be used by B (or other processes delegated by B) to return the result to A . Moreover, A waits for such a result, via an input at p . Therefore continuation names are *linear* [3] — they may only be used once — and *input receptive* [16] — the input-end of the name must be made available as soon as the name is created; and they are output-controlled: they carry the thread in output.

In short, the ‘well-bracketing’ type system defined in this section refines the type discipline for sequentiality by adding linear-receptive names and enforcing a stack discipline on the usage of such names. Proof techniques for well-bracketing will be studied in Section 4.4.

Thus, with well-bracketing, we have three kinds of names: output-controlled names (ranged over by $x, y, z, \dots$) and input-controlled names (ranged over by $u, v, w, \dots$), as in the previous section; and continuation names, ranged over by $p, q, r, \dots$. As before, names $a, b, c, \dots$ range over the union of output- and input-controlled names.

$$\begin{array}{c}
\text{WB-OUT1} \frac{}{p : \circ \vdash_{\text{wb}} \bar{p}\langle \tilde{a} \rangle} \quad \text{WB-OUT2} \frac{}{p : \circ \vdash_{\text{wb}} \bar{x}\langle \tilde{a}, p \rangle} \quad \text{WB-OUT3} \frac{}{\emptyset \vdash_{\text{wb}} \bar{u}\langle \tilde{a} \rangle} \\
\\
\text{WB-INP1} \frac{p : \circ \vdash_{\text{wb}} P \quad p \neq q}{q : \mathbf{i}, p : \circ \vdash_{\text{wb}} q(\tilde{a}).P} \quad \text{WB-INP2} \frac{p : \circ \vdash_{\text{wb}} P}{\emptyset \vdash_{\text{wb}} x(\tilde{a}, p).P, !x(\tilde{a}, p).P} \quad \text{WB-INP3} \frac{p : \circ \vdash_{\text{wb}} P}{p : \circ \vdash_{\text{wb}} u(\tilde{a}).P} \\
\\
\text{WB-NIL} \frac{}{\emptyset \vdash_{\text{wb}} \mathbf{0}} \quad \text{WB-RES1} \frac{\xi, p : \circ, p : \mathbf{i}, \sigma' \vdash_{\text{wb}} P}{\xi, \sigma' \vdash_{\text{wb}} (\nu p)P} \quad \text{WB-RES2} \frac{\sigma \vdash_{\text{wb}} P}{\sigma \vdash_{\text{wb}} (\nu p)P} \quad p \notin \sigma \quad \text{WB-RES3} \frac{\sigma \vdash_{\text{wb}} P}{\sigma \vdash_{\text{wb}} (\nu a)P} \\
\\
\text{WB-MAT} \frac{\emptyset \vdash_{\text{wb}} P}{\emptyset \vdash_{\text{wb}} [a = b]P} \quad \text{WB-PAR} \frac{\sigma \vdash_{\text{wb}} P \quad \sigma' \vdash_{\text{wb}} Q}{\sigma'' \vdash_{\text{wb}} P \mid Q} \quad \sigma'' \in \text{inter}(\sigma; \sigma') \quad \text{WB-TAU} \frac{\sigma \vdash_{\text{wb}} P}{\sigma \vdash_{\text{wb}} \tau.P} \quad |\sigma| \leq 1 \\
\\
\text{WB-SUM} \frac{\sigma \vdash_{\text{wb}} P \quad \sigma \vdash_{\text{wb}} Q}{\sigma \vdash_{\text{wb}} P + Q} \quad |\sigma| \leq 1
\end{array}$$

Fig. 5. Type system for well-bracketing.

Continuation names may only be sent at output-controlled names. Indeed, any output at an output-controlled name must carry exactly one continuation name. Without this constraint the type system for well-bracketing would be more complex, and it is unclear whether it would be useful in practice. Additionally, as a consequence of this constraint, sending a continuation means transmitting the thread, so that continuation names somehow represent the thread. By convention, we assume that, in a tuple of names transmitted over an output-controlled name, the last name is a continuation name. We write $\tilde{a}, p$ for such a tuple of names. The type system is presented in Fig. 5. Judgements are of the form

$$\sigma \vdash_{\text{wb}} P$$

where σ is a *stack*, namely a sequence of *input*- and *output*-tagged continuation names, in which the input and output tags alternate, always terminating with an output tag unless the sequence is empty:

$$\begin{array}{l}
\sigma ::= \sigma_0 \mid \sigma_1 \\
\sigma_0 ::= p : \circ, \sigma_1 \quad \sigma_1 ::= p : \mathbf{i}, \sigma_0 \mid \emptyset
\end{array}$$

Moreover: a name may appear at most once with a given tag; and, if a name appears with both tags, then the input occurrence should immediately follow the output occurrence, as for p in $p' : \mathbf{i}, p : \circ, p : \mathbf{i}, \sigma$. We write $p \in \sigma$ if name p appears in σ , and $|\sigma|$ for the length of the sequence σ .

Intuitively, a stack expresses the expected usage of the free continuation names in a process. For instance, if

$$p_1 : \circ, p_2 : \mathbf{i}, p_3 : \circ, p_3 : \mathbf{i}, p_4 : \circ \vdash_{\text{wb}} P$$

then $p_1, \dots, p_4$ are the free continuation names in P ; among these, p_1 will be used first, in an output (p_1 may be the subject or an object of the output); then p_2 will be used, in an input interaction with the environment. P possesses both the output and the input capability on p_3 , and may use both capabilities by performing a reduction at p_3 ; or P may transmit the output capability and then use the input one; the computation for P terminates with an output at p_4 . This behavior however concerns only the free continuation names of P : at any time when an output usage is expected, P may decide to create a new continuation name and send it out, maintaining its input end. The Subject Reduction Theorem 5 will formalize the behavior concerning continuation names in stacks. As simple examples of typing, we can derive

$$p : \circ \vdash_{\text{wb}} \bar{p}\langle a \rangle \quad \text{and} \quad p : \circ \vdash_{\text{wb}} u(a).\bar{p}\langle a \rangle$$

In the latter typing, by rule WB-INP3, an input at an input-controlled name has the thread, and does not affect the stack because u is not a continuation name.

The same stack can be used to type a process that invokes a service at x before sending the result at p , as in

$$p : \circ \vdash_{\text{wb}} (\nu q)(\bar{x}\langle b, q \rangle | q(c).\bar{p}\langle c \rangle)$$

where q is a fresh continuation name created when calling x . To type the process without the restriction at q , the stack should mention the input and output capabilities for q :

$$q : \circ, q : \mathbf{i}, p : \circ \vdash_{\text{wb}} \bar{x}\langle b, q \rangle | q(c).\bar{p}\langle c \rangle$$

For another example, the process

$$P_0 \stackrel{\text{def}}{=} p(a).\bar{p}'\langle a \rangle | q(b).q'\langle b \rangle$$

can be typed using two stacks: we have both $p : \mathbf{i}, p' : \circ, q : \mathbf{i}, q' : \circ \vdash_{\text{wb}} P_0$ and $q : \mathbf{i}, q' : \circ, p : \mathbf{i}, p' : \circ \vdash_{\text{wb}} P_0$. The choice of the stack depends on whether the call answered at p has been made before or after the call answered at q .

We comment on the rules of the type system. In **WB-OUT1** and **WB-OUT2** the obligation in the stack is fulfilled (the output capability on the only name in the stack is used in **WB-OUT1** and transmitted in **WB-OUT2**). Note, in **WB-OUT2**, that an output at an output-controlled name x , as in the system with sequentiality, passes the thread to another process (i.e., it is like calling a service x). As explained above, the last name in the tuple transmitted at x is a continuation name, and the only one being transmitted (to enforce the stack discipline). In contrast, in **WB-OUT3** an output at an input-controlled name does not own the thread and therefore may not carry continuation names.

In **WB-INP1** the input-tagged name on top of the stack is used. Rule **WB-INP2** is the complement of **WB-OUT2**. In **WB-INP3**, an input at an input-controlled name maintains the thread. In all rules for input and τ prefixes, the stack in the premise of the rules may not contain input-tagged continuation names because their input capability must be unguarded (as they are receptive names). The same occurs in rule **WB-SUM**, following [16] where choice on inputs at receptive names is disallowed (though the constraint could be relaxed). Matching is allowed on plain names, but not on continuation names; this is typical of type systems where the input and output capabilities on names are separate [1]; moreover, no continuation name may appear in the process underneath, to make sure that the obligations on continuation names are not eschewed. In **WB-RES1** a continuation name is created, and then its output and input capabilities are inserted into the stack. In the rule, ξ, σ is a decomposition of the stack for $\nu p P$ where σ is a stack beginning with an output tag; hence ξ is either empty or it is of the form $\sigma', p : i$, (i.e. ξ is an initial prefix of the stack, either empty or ending with an input tag). Being able to insert capabilities in the middle of the stack is useful for technical reasons. First, this guarantees commutativity of restrictions, i.e., if $(\nu p)(\nu q)P$ can be typed, then so can $(\nu q)(\nu p)P$. Second, it makes it possible to type processes like $\nu q (p.\bar{q}|q.p')$, where the output at q , the restricted name, must take place after the input at p . Rule **WB-RES2** is for continuation names that do not appear in the body of the restriction (this form of rule is common in type systems for linearity, to simplify the assertion of Subject Reduction). In rule **WB-PAR**, the typing stack is split to type the two process components P_1 and P_2 ; splitting of the typing is usual in type systems with linearity. Here, however, the split must respect the order of the names. That is, the stack in the conclusion should be an interleaving of the two stacks in the premises, as by the following definition.

Note also that while stacks are *finite* sequences of continuation names, a well-typed process may well engage in infinite computations.

Definition 8 (Interleaving).

We write $\sigma_1 \in \text{inter}(=2; =3)$ if (i) σ_1 is a stack, and (ii) σ_1 is an interleaving of σ_2 and σ_3 as by the following inductive rules:

1. $\emptyset \in \text{inter}(\emptyset; \emptyset)$
2. $p : \star, \sigma_1 \in \text{inter}(=2; p : \star, =3)$ if $\sigma_1 \in \text{inter}(=2; =3)$ for $\star \in \{o, i\}$
3. $\sigma_1 \in \text{inter}(=3; =2)$ if $\sigma_1 \in \text{inter}(=2; =3)$

If a name appears both in σ_2 and in σ_3 with the same tag, then $\text{inter}(=2; =3)$ do not contain any stack. Being stack-ordered means that $p.\bar{q}|q.\bar{p}$ cannot be typed. Indeed, the left process would require p before q in the stack, whereas the right process needs the opposite. In rule **RES1**, having the possibility to add names in the middle of the stack is mandatory to preserve typability after reduction. Consider for instance:

$$(\nu q)((\nu p)(\bar{b}(p)|p.\bar{q})|q.\bar{p}') \xrightarrow{(\nu p)\bar{b}(p)} (\nu q)(p.\bar{q}|q.\bar{p}')$$

To type the derivative of the transition above, we have to use rule **RES1** with $\xi = p : i$ and $\sigma' = p' : o$.

Typability in the type system of Fig. 5 implies typability in the type system for sequentiality. Indeed, when $\sigma \vdash_{\text{wb}} P$, if the first name in σ is output-tagged then P is active, otherwise P is inactive. We write seq for the function that ‘forgets’ the well-bracketing information in a stack, therefore $\text{seq}(\sigma) = 1$ if $\sigma = p : o, \sigma'$, for some p and σ' , and $\text{seq}(\sigma) = 0$ otherwise.

Proposition 6. If $\sigma \vdash_{\text{wb}} P$ then also $\text{seq}(\sigma) \vdash_s P$.

In Definition 9, we extend type-allowed transitions to processes with continuation names. As previously, we must ensure that the process is typed, and that the transition is allowed by sequentiality (clauses (1) and (2) below). Clause (3) says that the first continuation name observed must be on top of the stack, and that the input or output capability on a continuation name may not be exercised by the environment when both capabilities are owned by the process.

Definition 9. We write $\sigma \models_{\text{wb}} P \xrightarrow{\mu} P'$ when

1. $\sigma \vdash_{\text{wb}} P$
2. $\text{seq}(\sigma) \models_s P \xrightarrow{\mu} P'$ and
3. for any name p , whenever $p \in \text{fn}(\mu)$ and $p \in \sigma$, then either $\sigma = p : o, \sigma'$ or $\sigma = p : i, \sigma'$ for some σ' ; moreover, if $p \in \sigma'$, then p is not the subject of μ .

We exploit type-allowed transitions to define transitions with stacks, which make explicit the evolution of the stack.

Definition 10. We note $[\sigma; P] \xrightarrow{\mu} [\sigma'; P']$ when $\sigma \models_{\text{wb}} P \xrightarrow{\mu} P'$ and

1. if $\mu = (\nu \tilde{b})\bar{p}(\tilde{a})$, then $\sigma = p : o, \sigma'$ (the initial stack necessarily starts with an output at p);
2. if $\mu = p(\tilde{a})$, then $\sigma = p : i, \sigma'$ (the initial stack necessarily starts with an input at p);
3. if $\mu = (\nu \tilde{c}, p)\bar{x}(\tilde{b}, p)$, then $\sigma' = p : i, \sigma$ (an input at p is pushed on top of the stack: we wait for the call at x to be answered at p);

4. if $\mu = (\nu\tilde{c})\tilde{x}(\tilde{b}, p)$, then $\sigma = p : \circ, \sigma'$ (p is not fresh here; transmission of p means delegating the obligation to answer at p , therefore this obligation must be on top of the initial stack);
5. if $\mu = x(\tilde{b}, p)$, then $\sigma' = p : \circ, \sigma$ (an output at p is pushed on top of the stack: after x has been called, we must answer at p);
6. if $\mu = \tau$, then for $\sigma = p : \circ, p : \mathbf{i}, \sigma''$ and $p \notin \text{fn}(P')$, we have $\sigma' = \sigma''$, otherwise $\sigma' = \sigma$ (an internal communication occurring on a free continuation name p must have its two occurrences on top of the stack, and these are removed).
7. otherwise $\sigma' = \sigma$ (input-controlled names do not interact with the stack).

In cases (1), (4) (resp. (2)), we must have $\sigma = p : \circ, \sigma''$ (resp. $p : \mathbf{i}, \sigma''$) by definition of type-allowed transitions. In clauses (1) and (2), the action is an input or an output at a continuation name that must be on top of σ , and is then removed. In clause (3), the action extrudes a continuation name, and then, following the stack discipline, the process waits for an answer on that name. In clause (4), emitting a free continuation name amounts to passing the output capability on that name to the environment. Dually, in clause (5), receiving a continuation name imposes to use it in output. In clause (6), a τ transition may come from an interaction at a continuation name, in which case σ is modified. It can also come from an interaction at a restricted name or from an internal choice; in such cases, σ is unchanged. Finally, in clause (7), communications on input-controlled names do not interact with continuation names, leaving the stack identical.

Lemma 6. If $\sigma \vdash_{\text{wb}} P$ and $P \xrightarrow{\mu} P'$ with $\mu = (\nu\tilde{c})\tilde{x}(\tilde{b}, p)$ for some $x, \tilde{b}, \tilde{c}, p$ or $\mu = \bar{p}(\tilde{b})$ for some $p, \tilde{b}$, then $\sigma = p : \circ, \sigma'$ for some σ' .

Proof. From the weakening to sequential typing and Lemma 1, we must have that $\sigma = q : \circ$, so we just have to show that $p = q$. This follows from an induction on the transition. The main insight for the parallel composition is that if $\sigma_1 \in \text{inter}(p : \circ, \sigma_2; \Rightarrow_3)$ (or symmetrically), then $\sigma_1 = p : \circ, \sigma'_1$ for some σ'_1 . $\square$

Theorem 5 (Subject Reduction). If $\sigma \vdash_{\text{wb}} P$ and $[\sigma; P] \xrightarrow{\mu} [\sigma'; P']$ then $\sigma' \vdash_{\text{wb}} P'$.

Proof. The proof is similar to that of Theorem 2, following an induction on the transition relation. We highlight the changes related to the stack.

- Case ARES : We focus on the case with $\mu = \tau$, and the restriction is over a continuation name. The other cases are simpler, as there is no interaction between the action and the variation in the stack. The assumption can then be rewritten as $[\sigma; (\nu p)P] \xrightarrow{\tau} [\sigma'; (\nu p)P']$. The judgement $\sigma \vdash_{\text{wb}} (\nu p)P$ comes either from WB-RES1 or WB-RES2. The latter case follows by straightforward induction as the type does not change, i.e. $\sigma \vdash_{\text{wb}} P$. From WB-RES1, we have $\xi, p : \circ, p : \mathbf{i}, \sigma' \vdash_{\text{wb}} P$ with $\xi, \sigma' = \sigma$. By induction hypothesis, either we have $\xi, p : \circ, p : \mathbf{i}, \sigma' \vdash_{\text{wb}} P'$ or we can further decompose the stack to $\xi, p : \circ, p : \mathbf{i}, \sigma' = q : \circ, q : \mathbf{i}, \sigma''$ with $\sigma'' \vdash_{\text{wb}} P'$. In the first case, we can conclude as $\sigma \vdash_{\text{wb}} (\nu p)P'$ by WB-RES1. In the other case, if ξ is empty (in which case $p = q$), then $\sigma' = \sigma''$, and thus $\sigma' \vdash_{\text{wb}} (\nu p)P$ by using the other typing rule for the restriction, i.e. WB-RES2. Otherwise, if ξ is not empty, we conclude using WB-RES1.
- Case AOPEN, and the restriction is over a continuation name: The assumption can then be rewritten as $[\sigma; (\nu p)P] \xrightarrow{(\nu\tilde{c}, p)\tilde{a}(\tilde{b}, p)} [\sigma'; P']$. We note $\tilde{\sigma}$ the stack that types P . As $\tilde{\sigma} \models_{\text{wb}} P \xrightarrow{(\nu\tilde{c}, p)\tilde{a}(\tilde{b}, p)} P'$, by Lemma 6, the head of the stack $\tilde{\sigma}$ is $p : \circ$. Thus, $(\nu p)P$ is typed using WB-RES1, i.e. $\tilde{\sigma} = p : \circ, p : \mathbf{i}, \sigma$. By induction hypothesis, we have $p : \mathbf{i}, \sigma \vdash_{\text{wb}} P'$ (Case (3)). We then conclude as Case 4 imposes $\sigma' = p : \mathbf{i}, \sigma$.
- Case APAR : Compared to the sequentiality case, this case follows directly from the induction hypothesis, by using Lemma 6 and Definition 8 to obtain the updated type.
- Case ACOMM : The assumption for a communication on an output-controlled name a can be rewritten as $[\sigma; P_1 | P_2] \xrightarrow{\tau} [\sigma'; (\nu\tilde{c})(P'_1 | P'_2)]$, with $P_1 \xrightarrow{a(\tilde{b}, p)} P'_1$ and $P_2 \xrightarrow{(\nu\tilde{c})\tilde{a}(\tilde{b}, p)} P'_2$. For communications on continuation names, the case is similar, while for communications on an input-controlled name u , the stacks are left identical so the result follows easily. We call σ_1 and σ_2 the stacks for P_1 and P_2 respectively. By construction, $\sigma \in \text{inter}(\Rightarrow_1; \Rightarrow_2)$. There are two cases depending on whether the continuation name p is bound or not.

In the former, the restriction is of the form $(\nu\tilde{c}, p)$ both in the action and the resulting process. As we have $p \notin \sigma_1$ and $p \notin \text{fn}((\nu\tilde{c})\tilde{a}(\tilde{b}, p))$, both transitions are type-allowed (Definition 9). By induction hypothesis, we then have $p : \circ, \sigma_1 \vdash_{\text{wb}} P'_1$ and $p : \mathbf{i}, \sigma_2 \vdash_{\text{wb}} P'_2$. The stack $p : \circ, p : \mathbf{i}, \sigma$ forms a correct interleaving of $p : \circ, \sigma_1$ and $p : \mathbf{i}, \sigma_2$, so then we have $\sigma \vdash_{\text{wb}} (\nu\tilde{c}, p)(P'_1 | P'_2)$.

In the latter case, by Lemma 6, $\sigma_2 = p : \circ, \sigma'_2$. By induction hypothesis on both sides, we have $\sigma'_2 \vdash_{\text{wb}} P'_2$. As there exists an interleaving of σ_1 and σ_2 , if $p \in \sigma_1$, p must be on top, with $\sigma_1 = p : \mathbf{i}, \sigma'_1$. Either way, it means $\sigma_1 \models_{\text{wb}} P_1 \xrightarrow{a(\tilde{b}, p)} P'_1$, so by induction hypothesis $p : \circ, \sigma_1 \vdash_{\text{wb}} P'_1$. The stack σ is still an interleaving of the two new stacks $p : \circ, \sigma_1$ and σ'_2 , and the new restrictions only mention output-controlled names so $\sigma \vdash_{\text{wb}} (\nu\tilde{c})(P'_1 | P'_2)$.

$\square$

If a process owns both the input and the output capability on a continuation name p , then the environment may not use p . Semantically, this is the same as having a restriction on p in the process. It is therefore safe, in the definition of barbed bisimulation and observability, to assume that all such restrictions are syntactically present, i.e., there is a single occurrence of any free continuation name. We call *clean* such processes.

Definition 11. A stack σ is *clean* if no name appears in σ both output- and input-tagged. A process P is *clean* if $\sigma \vdash_{\text{wb}} P$ for some clean σ .

On clean processes, typing is preserved by reduction.

Proposition 7. *If $\sigma \vdash_{\text{wb}} P$ for σ clean, then $P \longrightarrow P'$ implies $\sigma \vdash_{\text{wb}} P'$.*

Defining barbed bisimulation on clean processes, we can use the seq function above to recast observability in the well-bracketing system from that in the sequentality system: thus, for σ clean, we have $\sigma \models_{\text{wb}} P \Downarrow_a$ (resp. $\sigma \models_{\text{wb}} P \Downarrow_p$) if $\text{seq}(\sigma) \models_s P \Downarrow_a$ (resp. $\text{seq}(\sigma) \models_s P \Downarrow_p$).

In the definition of barbed equivalence, the contexts testing the processes must be clean. Writing $\simeq^\sigma$ for barbed equivalence at σ , we have $P \simeq^\sigma Q$ if $\sigma \vdash_{\text{wb}} P, Q$, and for any clean σ' and any σ'/σ static context E , it holds that $E[P] \approx^{\sigma'} E[Q]$ (note that σ itself need not be clean).

4.2. Discreet processes

In this section we put forward the subclass of *discreet* processes, in which all continuation names that are exported must be private, and show how to transform any process into a discreet one. Then, on discreet processes:

- (1) we express a behavioral property that formalizes the stack-like discipline on the usage of continuation names;
- (2) we develop proof techniques, in form of labeled bisimilarities, to reason about the behavior of well-typed processes.

(Concerning (2), while the technical details are quite different, we follow the approach of proof techniques for receptive names in [16], where the techniques are first defined on processes where only fresh names may be sent.)

Definition 12 (Discreet processes). A process P is *discreet* if any free continuation name $p \in \text{fn}(P)$ may not appear in the object of an output, and, in any sub-process $x(\tilde{a}, q).Q$, the same holds for q in Q .

The definition is extended to contexts, yielding *discreet contexts*.

If E and P are discreet, then so is $E[P]$. We can transform all well-typed processes into discreet processes using the law in Lemma 7 below. The law transforms the output of a *global* continuation name p into the output of a *local* name q . In general, all outputs of continuation names in a process P are local, as a global output corresponds to P delegating a stack-like obligation to another process. In other words, in general the transformation of a non-discreet process into a discreet one will modify only a few outputs of the initial process. The law in Lemma 7 is valid for barbed congruence, not just barbed equivalence, and may therefore be applied to any component of a given process. The lemma is essentially an instance of the general law of asynchronous π -calculus about transforming the communication of a global name into the communication of a local name ‘linked’ to the global name, provided that only the output capability of those names is emitted [20].

Lemma 7. $\bar{x}(\tilde{a}, p) \cong^{p:\circ} (\nu q)(\bar{x}(\tilde{a}, q)|q(\tilde{b}).\bar{p}(\tilde{b}))$.

Thus, in the light of Lemma 7, in the definition of barbed equivalence (and congruence) it is sufficient to consider discreet contexts.

A discreet process may only export private continuation names. Dually, the process may only receive fresh continuation names from a discreet context. We call *discreet* the transitions that satisfy this property.

Definition 13 (Discreet transitions). A typed transition $\sigma \models_{\text{wb}} P \xrightarrow{\mu} P'$ is *discreet* if any continuation name in the object of μ is not free in σ (and hence also in P).

Lemma 8. *If P is discreet, and $\sigma \models_{\text{wb}} P \xrightarrow{\mu} P'$ is discreet, then P' is discreet. If, moreover, P is clean, then so is P' .*

4.3. The well-bracketing property on traces

Following game semantics [21], we formalize well-bracketing, that is, the stack-like behavior of continuation names for well-typed processes, using traces of actions. In this section, all processes are discreet and clean. A trace for such a process is obtained from a sequence of discreet transitions emanating from the process, with the expected freshness conditions to avoid ambiguity among names.

Definition 14 (Trace). A sequence of actions $\mu_1, \dots, \mu_n$ is a *trace for a (discreet and clean) process P_0 and a stack σ_0* if there are $\sigma_1, \dots, \sigma_n, P_1, \dots, P_n$ such that for all $0 \leq j < n$ we have $[\sigma_j; P_j] \xrightarrow{\mu_{j+1}} [\sigma_{j+1}; P_{j+1}]$, where the transition is discreet, and moreover all continuation names appearing as object in μ_{j+1} are fresh (i.e., the names may not appear in any μ_i for $i \leq j$).

The notion of discreet transition already imposes that continuation names in object position do not appear free in the process. The final condition in Definition 14 on continuation names ensures us that for actions like $(\nu p)\bar{x}(\tilde{a}, p)$, name p is fresh, and that after an action $\bar{p}(\tilde{a})$ (thus after the only allowed interaction at p has been played), name p cannot be reintroduced, e.g., in an action $x(\tilde{a}, p)$. We simply say that $\mu_1, \dots, \mu_n$ is a trace, or is a trace for P , when the stack or the process are clear from the context.

Lemma 9. *Suppose we have a trace $\mu_1, \dots, \mu_n$ with $(\sigma_j, P_j)_{0 \leq j < n}$ such that $[\sigma_j; P_j] \xrightarrow{\mu_{j+1}} [\sigma_{j+1}; P_{j+1}]$ for $0 \leq j < n$. If $\sigma_0 = \sigma_n$, then there exists some ρ_i (resp. ξ_i), such that $\sigma_i = \rho_i, \sigma_0$ (resp. ξ_i, σ_0) for all i .*

This lemma formalizes the idea that when looking at the evolution of the stacks for a given trace, names inside a stack may not reappear later on if they are removed. Thus, if at some point we obtain an identical stack, then for all the intermediate stacks, we should find the original stack as a substack.

The well-bracketing property is best described with the notion of questions and answers.

Definition 15. For a trace $\mu_1, \dots, \mu_n$, we set $\mu_i \curvearrowright \mu_j$ if $i < j$ and:

1. either $\mu_i = (\nu \tilde{c}, p) \tilde{a}(\tilde{b}, p)$ and $\mu_j = p(\tilde{a}')$,
2. or $\mu_i = a(\tilde{b}, p)$ and $\mu_j = (\nu \tilde{c}) \tilde{p}(\tilde{a}')$.

Actions μ_i (with a continuation name in object position) are called *questions*, while actions μ_j (with a continuation name in subject position) are called *answers*.

A discreet transition is either an internal transition, or a communication on an input-controlled name, or a question, or an answer. A question mentioning a continuation name p is matched by an answer at the same name p . When questions and answers are seen as delimiters ($[\cdot]_p, [\cdot]_p$, different for each continuation name), a well-bracketed trace is a substring of a Dyck word. It can be expressed as follows.

Remark 1. For a discreet transition $[\sigma; P] \xrightarrow{\mu} [\sigma'; P']$, the value $|\sigma'| - |\sigma|$ is 1 for a question, 0 for an internal action or a communication on an input-controlled name, and -1 for an answer.

Lemma 10 (Uniqueness). *Given a trace $\mu_1, \dots, \mu_n$, if $\mu_i \curvearrowright \mu_j$ and $\mu_{i'} \curvearrowright \mu_{j'}$, then we have $(i = i' \text{ iff } j = j')$.*

Definition 16 (Well-bracketing). A trace $\mu_1, \dots, \mu_n$ is *well-bracketed* if for all $i < j$, if μ_i is a question and μ_j is an answer with $\mu_i \curvearrowright \mu_j$ and $\mu_k \not\curvearrowright \mu_j$ for all $i < k < j$, then $\mu_i \curvearrowright \mu_j$.

To prove that all traces are well-bracketed, (Proposition 8), we need the following property relating questions and answers to stacks.

Lemma 11. *Let $\mu_1, \dots, \mu_n$ be a trace, for $n \geq 2$, and $\sigma_0, \dots, \sigma_n$ be the corresponding stacks, as in Definition 14. Suppose $\sigma_0 = \sigma_n$, and for all i with $0 < i < n$, we have $|\sigma_i| > |\sigma_0|$. Then μ_1 is a question, μ_n is an answer, and $\mu_1 \curvearrowright \mu_n$.*

Proposition 8. *Any trace (as by Definition 14) is well-bracketed.*

4.4. Bisimulation and full abstraction

As in Section 3, a *wb-typed relation on processes* is a set of triplets (σ, P, Q) with $\sigma \vdash_{\text{wb}} P, Q$.

Definition 17 (WB-Bisimulation). A *wb-typed relation* $\mathcal{R}$ on discreet processes is a *wb-bisimulation* if whenever $(\sigma, P, Q) \in \mathcal{R}$ and $[\sigma; P] \xrightarrow{\mu} [\sigma'; P']$ is discreet, then one of the three following clauses holds:

1. there is Q' with $Q \xRightarrow{\mu} Q'$ and $(\sigma', P', Q') \in \mathcal{R}$
2. $\mu = x(\tilde{a}, p)$ and for some fresh q , there is Q' with $Q | (\nu q)(\tilde{x}(\tilde{a}, q) | q(\tilde{b}). \tilde{p}(\tilde{b})) \Rightarrow Q'$ and $(\sigma', P', Q') \in \mathcal{R}$
3. $\mu = u(\tilde{a})$ and there is Q' with $Q | \tilde{u}(\tilde{a}) \Rightarrow Q'$ and $(\sigma', P', Q') \in \mathcal{R}$,

and symmetrically for the transitions from Q .

Processes P and Q are *wb-bisimilar at σ* , noted $P \approx_{\text{wb}}^{\sigma} Q$, if $(\sigma, P, Q) \in \mathcal{R}$ for some *wb-bisimulation* $\mathcal{R}$.

Compared to Definition 6, the clause for input actions is here split into two clauses. In clause (2), we apply Lemma 7 to obtain a discreet process. The new bisimilarity is sound with respect to barbed equivalence for all discreet processes. For this, the main result concerns preservation by parallel composition. Before proving this result, we show that *up-to restrictions* is a sound proof technique. For a typed relation $\mathcal{R}$, we write $(\sigma, P, Q) \in \mathcal{R}^v$ when there exists σ' , a σ/σ' context E of the form $(\nu \tilde{c})([\cdot])$, and processes P', Q' such that $Q \equiv E[Q']$, $P \equiv E[P']$ and $(\sigma', P', Q') \in \mathcal{R}$. This gives rise to *wb-bisimulation up to restrictions*, which a weaker version of Definition 18, using $\mathcal{R}^v$ instead of $\mathcal{R}^C$.

Lemma 12. *If $\mathcal{R}$ is a *wb-bisimulation up-to restrictions*, then $(\sigma, P, Q) \in \mathcal{R}$ implies $P \approx_{\text{wb}}^{\sigma} Q$.*

Proof. The proof amounts to proving that $\mathcal{R}^v$ is a *wb-bisimulation*. Take $(\sigma, (\nu \tilde{c})P, (\nu \tilde{c})Q) \in \mathcal{R}^v$, and a discreet transition $[\sigma; (\nu \tilde{c})P] \xrightarrow{\mu} [\sigma_2; P']$. We know $(\sigma', P, Q) \in \mathcal{R}$ for some σ' , so by a case analysis, we have the following cases:

- $P \xrightarrow{\mu} P''$ and $P' \equiv (\nu \tilde{c})P''$;
- μ is of the form $(\nu \tilde{c}_2, \tilde{c}_3, p) \tilde{a}(\tilde{b}, p)$ with $\tilde{c} = \tilde{c}_1, \tilde{c}_2$, and $P \xrightarrow{(\nu \tilde{c}_3, p) \tilde{a}(\tilde{b}, p)} P''$ and $P' \equiv (\nu \tilde{c}_1, \tilde{c}_3)P''$.

Note that by definition of $\mathcal{R}$, the process P is discreet, so in the second case, the continuation name p must be bound. In both cases, we have $\sigma \models_{\text{wb}} P \xrightarrow{\mu'} P''$ for some μ' and P'' such that $P' \equiv (\nu \tilde{d})P''$ for some $\tilde{d}$, and that transition is discreet. By Definition 10, there exists σ'_2 such that $[\sigma'; P] \xrightarrow{\mu'} [\sigma'_2; P'']$. By subject reduction (Theorem 5), we have $\sigma'_2 \vdash_{\text{wb}} P''$, and $\sigma_2 \vdash_{\text{wb}} P'$ with $(\nu \tilde{d})([\cdot])$ being a σ_2/σ'_2 context.

As $\mathcal{R}$ is a *wb-bisimulation up-to restrictions*, the typed transition entails that, e.g. there exists Q' with $Q \xRightarrow{\mu'} Q''$, and $(\sigma'_2, P'', Q'') \in \mathcal{R}^v$: we focus on clause (1) of Definition 18, the other two clauses are handled similarly. For the transition of Q , we have $(\nu \tilde{c})Q \xRightarrow{\mu} (\nu \tilde{d})Q''$. As $(\sigma_2, (\nu \tilde{d})P'', (\nu \tilde{d})Q'') \in (\mathcal{R}^v)^v$, we can conclude as $(\mathcal{R}^v)^v = \mathcal{R}^v$. $\square$

In [Lemma 13](#) below, we note that even if P, R are clean, $P|R$ needs not be so.

Lemma 13 (Parallel composition). *If $P \approx_{\text{wb}}^{\sigma} Q$, then for any discreet process R and stacks σ', σ'' such that $\sigma' \vdash_{\text{wb}} R$ and $\sigma'' \in \text{inter}(=; =')$, we have $P|R \approx_{\text{wb}}^{\sigma''} Q|R$.*

Proof. We show that the typed relation

$$\{(\sigma'', P|R, Q|R) \text{ s.t. } P \approx_{\text{wb}}^{\sigma} Q, \sigma' \vdash_{\text{wb}} R, R \text{ discreet}, \sigma'' \in \text{inter}(=; =')\}$$

is a wb-bisimulation up-to restrictions. We call $\mathcal{R}$ this typed relation. Take $(\sigma'', P|R, Q|R) \in \mathcal{R}$, and a discreet transition $[\sigma''; P|R] \xrightarrow{\mu} [\sigma_2''; P_2]$, we have three cases:

- $P_2 \equiv P'|R$ (rule **PAR**): By [Definition 10](#), there exists a σ_2 such that $[\sigma; P] \xrightarrow{\mu} [\sigma_2; P']$, and $\sigma_2'' \in \text{inter}(=; =')$. As $P \approx_{\text{wb}}^{\sigma} Q$, we have one of the following:
 - Either there exists Q' such that $Q \xRightarrow{\hat{\mu}} Q'$ and $P' \approx_{\text{wb}}^{\sigma_2} Q'$. This means we have $\sigma_2 \vdash_{\text{wb}} Q'$ and thus $\sigma_2'' \vdash_{\text{wb}} Q'|R$, so $(\sigma_2'', P'|R, Q'|R) \in \mathcal{R} \subseteq \mathcal{R}^{\vee}$. We then conclude as $Q|R \xRightarrow{\hat{\mu}} Q'|R$ by rule **PAR**.
 - Or $\mu = x(\tilde{a}, p)$ and for some fresh q , $Q|(\nu q)(\tilde{x}(\tilde{a}, q)|q(\tilde{b}).\tilde{p}(\tilde{b})) \Rightarrow Q'$ and $P' \approx_{\text{wb}}^{\sigma_2} Q'$. Thus, using the same reasoning as above, we have $Q|R|(\nu q)(\tilde{x}(\tilde{a}, q)|q(\tilde{b}).\tilde{p}(\tilde{b})) \Rightarrow Q'|R$ and $(\sigma_2'', P'|R, Q'|R) \in \mathcal{R} \subseteq \mathcal{R}^{\vee}$.
 - Or $\mu = u(\tilde{a})$ and $Q|\tilde{u}(\tilde{a}) \Rightarrow Q'$ and $P' \approx_{\text{wb}}^{\sigma_2} Q'$ which is dealt identically to the second case.
- $P_2 \equiv P|R'$ (rule **PAR**): Similarly to the previous case, there exists σ_2' such that $[\sigma'; R] \xrightarrow{\mu} [\sigma_2'; R']$, and $\sigma_2'' \in \text{inter}(=; =')$. We can then conclude directly as $Q|R \xrightarrow{\mu} Q|R'$ and $\sigma_2'' \vdash_{\text{wb}} Q|R'$ so $(\sigma_2'', P|R', Q|R') \in \mathcal{R}$.
- $P_2 \equiv (\nu \tilde{c}, p)(P'|R')$ (rule **COMM**): As for the proof of [Theorem 3](#), we focus on the case where $P \xrightarrow{(\nu \tilde{c}, p)\tilde{a}(\tilde{b}, p)} P'$ and $R \xrightarrow{a(\tilde{b}, p)} R'$. Similarly to the proof of [Theorem 5](#), we have $[\sigma; P] \xrightarrow{(\nu \tilde{c}, p)\tilde{a}(\tilde{b}, p)} [p : \mathfrak{i}, \sigma; P']$ and $[\sigma'; R] \xrightarrow{a(\tilde{b}, p)} [p : \circ, \sigma'; R']$. Thus there exists Q' such that $Q \xRightarrow{\hat{\mu}} Q'$ and $P' \approx_{\text{wb}}^{p: \mathfrak{i}, \sigma} Q'$. This implies that both P' and Q' can be typed by $p : \mathfrak{i}, \sigma$, and thus that $\sigma_2'' \vdash_{\text{wb}} (\nu \tilde{c}, p)(Q'|R')$. By looking at the type derivation, we find a stack σ_3 that types both $P'|R'$ and $Q'|R'$, and as such, we have $(\sigma_3, P'|R', Q'|R') \in \mathcal{R}$. By construction of σ_3 and $\mathcal{R}^{\vee}$, it leads to $(\sigma_2'', (\nu \tilde{c}, p)(P'|R'), (\nu \tilde{c}, p)(Q'|R')) \in \mathcal{R}^{\vee}$. In combination with the fact that $Q|R \Rightarrow (\nu \tilde{c}, p)(Q'|R')$, we can conclude.

Thus $\mathcal{R}$ is a wb-bisimulation up-to restrictions and by [Lemma 12](#), $\mathcal{R}$ is included in wb-bisimilarity. $\square$

[Lemma 13](#) entails soundness for $\approx_{\text{wb}}^{\sigma}$ under all discreet static contexts. By [Lemma 7](#), we can then replace any non-discreet context with a discreet one, without changing its behavior. We can thus derive (full) soundness.

Theorem 6 (Soundness). $\approx_{\text{wb}}^{\sigma} \subseteq \simeq^{\sigma}$.

Now we further refine the coinductive technique given by $\approx_{\text{wb}}^{\sigma}$ by introducing some up-to techniques, which make it possible to work with smaller relations. We write $P \rightarrow_d P'$ when the reduction $P \rightarrow P'$ is *deterministic*, meaning that whenever $P \xrightarrow{\mu} P''$, then $\mu = \tau$ and $P' \equiv P''$. Similarly, we write $P \Rightarrow_d P'$ if all reduction steps are deterministic. Moreover, for a relation $\mathcal{R}$, we write $(\sigma, P, Q) \in \Rightarrow_d \mathcal{R}^C$ when there exists a stack σ' , a σ/σ' context E , and processes P', Q' such that $Q \equiv E[Q']$, $P \Rightarrow_d E[P']$ and $(\sigma', P', Q') \in \mathcal{R}$.

Definition 18 (Up-to static contexts and up-to deterministic reductions). A wb-typed relation $\mathcal{R}$ on discreet processes is a *wb-bisimulation up-to static contexts and up-to deterministic reductions* if whenever $(\sigma, P, Q) \in \mathcal{R}$, for any discreet transition $[\sigma; P] \xrightarrow{\mu} [\sigma'; P']$, one of the following clauses holds:

1. there is Q' with $Q \xRightarrow{\hat{\mu}} Q'$ and $(\sigma', P', Q') \in \Rightarrow_d \mathcal{R}^C$,
2. $\mu = x(\tilde{a}, p)$ and for some fresh q , there is Q' with $Q|(\nu q)(\tilde{x}(\tilde{a}, q)|q(\tilde{b}).\tilde{p}(\tilde{b})) \Rightarrow Q'$ and $(\sigma', P', Q') \in \Rightarrow_d \mathcal{R}^C$,
3. $\mu = u(\tilde{a})$ and, there is Q' with $Q|\tilde{u}(\tilde{a}) \Rightarrow Q'$ and $(\sigma', P', Q') \in \Rightarrow_d \mathcal{R}^C$,

and symmetrically for the transitions from Q .

Remark 2. In the definition of $\Rightarrow_d \mathcal{R}^C$, the reduction $P \Rightarrow_d E[P']$ cannot change the typing. Thus, interactions on free continuation names, as in $\tilde{p}|p.Q \rightarrow_d Q$, may not be involved.

Lemma 14. *If $\mathcal{R}$ is a wb-bisimulation up-to static contexts and up-to deterministic reductions, then $(\sigma, P, Q) \in \mathcal{R}$ implies $P \approx_{\text{wb}}^{\sigma} Q$.*

4.4.1. Completeness

As in [Section 3](#), we prove completeness for processes that only use output-controlled names.

Theorem 7 (Completeness). *For all image-finite, discreet and clean processes P, Q that only use output-controlled names, and for all σ , if $P \simeq^\sigma Q$ then $P \approx_{\text{wb}}^\sigma Q$.*

As for [Theorem 4](#), the crux of the proof is defining the discriminating static contexts — the counterpart of [Proposition 5](#). The additional difficulty is related to receptiveness of continuation names: we cannot use $z.R$ or $G_R + T$, as in [Section 3](#), when the tester process, R or G_R , contains an input at a free continuation name.

Suppose $\sigma \vdash_{\text{wb}} P$, for P discreet and clean. We decompose σ as $\xi, p_1 : \circ, q_1 : \mathbf{i}, \dots, p_{n-1} : \circ, q_{n-1} : \mathbf{i}, p_n : \circ$ for $\xi = \emptyset$ or $\xi = q : \mathbf{i}$, and then define, for fresh q_n and $\tilde{x}_i$,

$$E_\sigma^{\tilde{x}_i} \stackrel{\text{def}}{=} (\nu \tilde{p}_i, \tilde{q}_i)(\llbracket \cdot \rrbracket \prod_{i \leq n} p_i(\tilde{y}).\tilde{x}_i(\tilde{y}, q_i)).$$

We have $\xi, q_n : \circ \vdash_{\text{wb}} E_\sigma^{\tilde{x}_i}[P]$. Intuitively, $E_\sigma^{\tilde{x}_i}$ forwards information from the p_i 's (which are in σ) to the x_i 's. Accordingly, the tester process can use names in $\tilde{x}_i$ (rather than in $\tilde{p}_i$), and can use them in input.

Let $\text{fn}_0(-)$ denote the set of free output-controlled names. We distinguish two cases. If $\xi = \emptyset$, then P is active. To follow the reasoning in the proof of [Theorem 4](#), we work with E of the form

$$(\nu \tilde{x})(E_\sigma^{\tilde{x}_i}(R + \sum_{y \in S} y(\tilde{y}', p).\bar{z}\langle p \rangle)),$$

for some set $S \supseteq \text{fn}_0(P) \cup \text{fn}_0(Q) \cup \tilde{x}_i$ and fresh z .

If $\xi = q : \mathbf{i}$, then P is inactive. We reason with E of the form $(\nu \tilde{x}, q)(E_\sigma^{\tilde{x}_i}[\bar{z}\langle q \rangle]z(q').R)$. By typing, only the continuation name q' received at z may appear free in R . Such q' will be instantiated with q and then R will use it to test the input at q from the tested processes. (A restriction on q is needed, as the overall process has to be clean.)

In both cases, the resulting E is a $q_n : \circ/\sigma$ static context where q_n is a fresh continuation name. More details on the proof, including the precise statement and proof of the counterpart of [Proposition 5](#), may be found in [Appendix B](#).

4.4.2. An example

We explain how the techniques we have introduced allow us to reason about a well-known example, the *well-bracketed state change* (sometimes called ‘awkward’, or ‘very awkward’, example) [22–24]. It is usually presented in ML thus:

$$\begin{aligned} M_1 &\stackrel{\text{def}}{=} \text{let } \ell = \text{ref0infuny} - > \\ &\quad (\ell := 0; y(); \ell := 1; y(); !\ell) \\ M_2 &\stackrel{\text{def}}{=} \text{funy} - > (y(); y(); 1) \end{aligned}$$

Function M_2 makes two calls to an external function y and returns 1. The other term, M_1 , between the two calls, modifies a local reference ℓ , which is then used to return the final result. Intuitively, equivalence between the two functions holds because: (i) the reference ℓ in M_1 represents a local state, not accessible from an external function; (ii) computation respects well-bracketing (e.g., the language does not have control operators like call/cc with which it would be possible to reinstall the value 0 at ℓ , and obtain 0 as final result).

Below are the translations of M_1 and M_2 , following a standard encoding of functions and references in $\Lambda\pi$, and using the notations for references from [Section 3.3.2](#):

$$\begin{aligned} \llbracket M_1 \rrbracket_{p'} &\stackrel{\text{def}}{=} (\nu x, \ell)(\bar{\ell}\langle 0 \rangle | Q) \quad \text{with} \\ Q &\stackrel{\text{def}}{=} \bar{p}'\langle x \rangle | !x(y, p).\text{wr}_\ell\langle 0 \rangle.(\nu q)(\bar{y}\langle q \rangle | \\ &\quad q.\text{wr}_\ell\langle 1 \rangle.(\nu r)(\bar{y}\langle r \rangle | r.\text{re}_\ell(n).\bar{p}\langle n \rangle)) \\ \llbracket M_2 \rrbracket_{p'} &\stackrel{\text{def}}{=} (\nu x)(\bar{p}'\langle x \rangle | !x(y, p).(\nu q)(\bar{y}\langle q \rangle | \\ &\quad q.(\nu r)(\bar{y}\langle r \rangle | r.\bar{p}\langle 1 \rangle))) \end{aligned}$$

$\llbracket M_1 \rrbracket_{p'}$ has a unique transition, $\llbracket M_1 \rrbracket_{p'} \xrightarrow{(\nu x)\bar{p}'\langle x \rangle} P_1$. Similarly, let P_2 be the unique derivative from $\llbracket M_2 \rrbracket_{p'}$. The equivalence between $\llbracket M_1 \rrbracket_{p'}$ and $\llbracket M_2 \rrbracket_{p'}$ follows immediately from $P_1 \approx_{\text{wb}}^\emptyset P_2$. To prove the latter, we exhibit a relation $\mathcal{R}$ containing the triple $(\emptyset, P_1, P_2)$ and show that $\mathcal{R}$ is a wb-bisimulation up-to deterministic reductions and static context. To highlight why the well-bracketing discipline is essential in the proof, consider the following triple that appears in $\mathcal{R}$:

$$\begin{aligned} &(\quad q_2 : \mathbf{i}, p_2 : \circ, r_1 : \mathbf{i}, p_1 : \circ, \\ &(\nu \ell)(\bar{\ell}\langle 0 \rangle | Q | q_2.\text{wr}_\ell\langle 1 \rangle.(\nu r_2)(\bar{y}\langle r_2 \rangle \\ &\quad | r_2.\text{re}_\ell(n).\bar{p}_2\langle n \rangle)) | r_1.\text{re}_\ell(n).\bar{p}_1\langle n \rangle), \\ &P_2 | q_2.(\nu r_2)(\bar{y}\langle r_2 \rangle | r_2.\bar{p}_2\langle 1 \rangle) | r_1.\bar{p}_1\langle 1 \rangle) \end{aligned}$$

Without the well-bracketing constraint, the first process in the triple could perform an input at r_1 , an internal transition, and finally an output $\bar{p}_1\langle 0 \rangle$. The second process cannot emit 0, which would allow us to distinguish P_1 and P_2 . With well-bracketing, since r_1 is not on top of the stack in the triple, the initial transition on r_1 is ruled out.

The details of the definition of $\mathcal{R}$ can be found in [Appendix C](#). In the same Appendix, we also discuss a simplified example, which exposes the main difficulties. The primary simplification consists in using linear functions, so that the $\lambda\pi$ terms do not need replication and can be finite. Some twisting in the ML terms is however necessary, as M_1 and M_2 become equivalent — even dropping well-bracketing — if they can be used at most once.

5. Related work and conclusions

Sequentiality is a form of linearity, hence the rules for our type system has similarities to, and borrow ideas from, type systems with linear types, in languages for concurrency or functional languages, including types for managing locks as in [\[8\]](#). The type system in [\[25\]](#) ensures that terms of the Ambient calculus are single-threaded, a notion similar to the sequentiality for $\lambda\pi$ examined in this paper. The type system in [\[26\]](#) has been designed so to make the encoding of PCF into the π -calculus fully abstract. The system therefore goes beyond sequentiality as described in our paper. For instance, the system presents a form of duality on types and ensures that computations are stateless, hence also deterministic. Indeed, the only behaviors inhabited by the types are those in the image of the PCF terms. Types ensure the uniqueness of the computation thread, and such a thread is carried by outputs (the thread cannot be carried by input processes, as in our system). The system [\[26\]](#) has been further refined in [\[10\]](#), adding causality information and acyclicity constraints, so to ensure strong normalization of well-typed processes. The issue of finding labeled bisimilarity characterisations of barbed equivalence or reduction-closed barbed equivalence is extensively discussed in [\[1\]](#); see also [\[27\]](#) for an example involving types.

Type systems for linearity and receptiveness in the π -calculus have been introduced in [\[3,16,28\]](#). The way we formulate well-bracketing ([Definition 16](#)) is inspired by ‘well-bracketed strategies’ in game semantics [\[21,29\]](#), used in functional programming languages and extensions thereof (they have in turn inspired type systems for π -calculi with stack-like information and input/output alternation, e.g., [\[26,30\]](#)). The notion of *well-bracketed control flow* is studied in the field of secure compilation, for a wider class of languages. In works like [\[31,32\]](#), the technique of fully abstract compilation guarantees control flow correctness (and, in particular, well-bracketing) against low-level attacks.

Several methods have been proposed to establish contextual equivalence of sequential programs that include higher-order and stateful computation, including the above-mentioned game semantics, (step-indexed Kripke) logical relations [\[22,23\]](#), dedicated forms of bisimulations designed on top of an operational semantics of the languages [\[33–37\]](#). Works like [\[24\]](#) or algorithmic game semantics [\[38\]](#), aim at automatically establishing contextual equivalences, by relying on model-checking techniques.

The main goal of this paper was to tailor some of the most prominent proof techniques in the π -calculus — those based on labeled bisimilarity — to the sequentiality and well-bracketing disciplines. This is instrumental to the use of the π -calculus as a model of programming languages, as sequentiality and well-bracketing are often found in programming languages or subsets of them. We have shown the usefulness of our techniques on a number of examples, that have mainly to do with the representation of functions and store — none of the equalities in the examples is valid in the ordinary bisimilarity of the calculus.

In [Section 3.3](#) we have combined our proof technique for sequentiality with techniques concerning the representation of references in π -calculus from [\[17\]](#). The resulting technique allows us in some cases to reason about programs with store without an explicit representation of the store (as usually required in the techniques in the literature, recalled above). This avoids universal quantifications on the possible values contained in the store, thus reducing the size of the relation to consider, sometimes making it finite. Further possibilities of reducing the size of relations may be possible by adapting the rich theory of the ‘up-to techniques’ for the ordinary π -calculus to our bisimilarities, as exemplified by the up-to technique considered in [Definition 18](#) and applied in [Section 4.4.2](#). By the time the writing of this paper has been completed, the ideas presented in this paper have been used to study behavioral equivalences for processes in which references may only contain first-order values [\[39,40\]](#).

Our treatment of sequentiality raises a few technical questions that deserve further investigation. We would like to see whether our proof of completeness ([Theorem 4](#)) could be extended to handle input-controlled names. Similarly, we do not know whether the result still holds if internal choice is disallowed in inactive processes. The usual encoding of an internal choice $\tau.P + \tau.Q$ in terms of parallel composition as $(\nu c)(\bar{c}|c.P|c.Q)$, for some fresh c , is not applicable because the latter process is active (for instance, the encoding is not valid within a context testing active processes). Indeed, if the result still holds, the current completeness proof would probably require some significant modifications. For similar reasons, it is unclear if and how our completeness proof could be tuned to handle reduction-closed variants of barbed equivalence [\[1,12\]](#).

In the *asynchronous* π -calculus considered in this paper, an interaction involves only one prefixed process (the input). Therefore, in the type systems, this process always acquires the control on the thread after the interaction. In a *synchronous* setting, in contrast, an interaction involves also an output prefix. Hence the type systems could be richer, specifying, for each name, where the control on the thread goes after an interaction at that name. The representation of references in [Section 3.3](#), however, might have to be revisited as it relies on the asynchronous model.

We have studied type systems and proof methods for sequentiality and well-bracketing in the pure π -calculus. We would like to experiment such techniques in a richer setting, or in more concrete applications. Good targets would be some of the many techniques and languages studied by Gian-Luigi Ferrari, notably: techniques for resource usage analysis and verification in languages where resources can be dynamically created, passed, and returned as results of calls (e.g., [\[41–43\]](#)); programming languages for distributed systems such as Klaim [\[44\]](#), where both processes and data can move, and types permit checking access rights violations of mobile agents relative to the localities with which the agents interact or where they intend migrating to; compiler optimisations; validation of transformations under various memory models.

CRedit authorship contribution statement

Daniel Hirschhoff: Writing – original draft, Conceptualization; **Enguerrand Prebet:** Writing – original draft, Methodology, Conceptualization; **Davide Sangiorgi:** Writing – original draft, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank the reviewers for a number of useful suggestions.

The third author would like to pay a special tribute to Gian-Luigi (“Giangi”) Ferrari, to whom this paper is dedicated. Giangi is, and has always been, a great scientist, with a culture both broad and deep — proof of this is the remarkable range of topics to which he has contributed. Indeed, having Giangi in a national or EU project has always been, scientifically, a pleasure, as well as a reassurance with the project reviewers. A spectacular example has been the EU project Profundis: how could we have coped with the reviewers without all Giangi’s work? On top of all this, Giangi is, and has always been, an exceptional person: always ready to share, to listen, to offer advice, and quietly and unfailingly attentive to the people around him.

Data availability

No data was used for the research described in the article.

Appendix A. Completeness for sequential bisimilarity

We recall [Proposition 5](#) before giving the proof.

Proof. We reason by induction on n . For $n = 0$, there is nothing to prove. For $n > 0$, suppose that $P \not\approx_s^{\eta,n} Q$. We may therefore suppose that there exists μ such that $[\eta; P] \xrightarrow{\mu} [\eta'; P']$ and for all Q' with $Q \xRightarrow{\hat{\mu}} Q'$ (or when $\mu = x(\tilde{y})$, all Q' with $Q|\bar{x}(\tilde{y}) \Rightarrow Q'$), we have $P' \not\approx_s^{\eta',n-1} Q'$.

We note $\{Q_i\}$, for $i \in I$, the set of all such Q' . This set is finite by hypothesis. We also write $S'_I = \bigcup_{i \in I} \text{fn}(Q_i) \cup \text{fn}(P')$, which is also a finite set.

We show that clause 1) holds (in the symmetrical case, where we consider a transition from $[\eta; Q]$, we establish the second clause). We first consider the case where P' is singular. The reasoning is simpler in this case.

First case: P' is singular.

Since P' is singular, $\eta' = 1$. For all i , as $P' \not\approx_s^1 Q_i$, we have $Q_i \Downarrow_{\bar{z}_i}$ for some z_i . By definition, S'_I contains all such z_i ’s. We consider the transition from P to P' . By [Theorem 2](#) (Subject Reduction), since P' is singular, P' is active, so μ cannot be an output. Hence there are two possible cases:

- when $\mu = \tau$, we can take $R \stackrel{\text{def}}{=} \mathbf{0}$.

We then have to show that for any $z, \tilde{x}, S, Q'$ with $Q \Rightarrow Q'$, taking $\bar{R} = \mathbf{0} + \sum_{y \in S} y(\tilde{y}).\bar{z}$ (since $\eta = 1$) and $E_0 = (\nu \tilde{x})([\cdot]|\bar{R})$, we have $E_0[P] \dot{\approx}^1 E_0[Q']$.

We reason by contradiction. We observe that $E_0[P] \rightarrow E_0[P']$, which is singular. Thus there exists a process T such that $E_0[Q'] \Rightarrow T$ and $E_0[P'] \dot{\approx}^1 T$, which entails, by [Lemma 4](#), that T is singular. Since $E_0[Q'] \Rightarrow T$, T can be of two forms, up to $\equiv$:

- $E_0[Q_i]$ for some $i \in I$ (with $Q' \Rightarrow Q_i$);
- $(\nu \tilde{x})(Q''|\bar{z})$.

Since T is singular, the second case is impossible, because by hypothesis $z \notin \tilde{x}$ so we would have $T \Downarrow_z$.

For the first case, notice that $\text{fn}(Q_i) \subseteq \text{fn}(Q)$ and $\text{fn}(P') \subseteq \text{fn}(P)$. Thus, $S'_I \subseteq S$. As $Q_i \Downarrow_{\bar{z}_i}$ and $z_i \in S'_I$; this means that $T \Downarrow_{\bar{z}}$, which contradicts the fact that T is singular.

- when $\mu = x'(\tilde{y})$, given a fresh name z' , we define

$$R \stackrel{\text{def}}{=} \bar{x}'(\tilde{y}) \mid \sum_{y \in S'_I} y(\tilde{y}).\bar{z}'.$$

We then have to show that for any $z, \tilde{x}, S, Q'$, with $Q \Rightarrow Q'$,

$$E_0[P] \dot{\approx}^1 E_0[Q'],$$

with $E_0 = (\nu \tilde{x})([\cdot]|\bar{z}.R|\bar{z})$, since in this case $\eta = 0$ (by [Theorem 2](#)).

We reason by contradiction and note $E_1 \stackrel{def}{=} (\nu \tilde{x})([\cdot] \mid \sum_{y \in S'_I} y(\tilde{y}'). \bar{z}')$.

We have $E_0[P] \rightarrow E_1[P']$ which is a singular process. Thus there should exist T such that $E_0[Q'] \Rightarrow T$ and $E_1[P'] \stackrel{\cdot}{\approx}^1 T$. We observe that T can be of 3 forms:

- $E_0[Q'']$ (with $Q' \Rightarrow Q''$)
- $E_1[Q_i]$ for some $i \in I$ (with $Q' \mid \bar{x}'(\tilde{y}') \Rightarrow Q_i$)
- $(\nu \tilde{x})(Q'' \mid \bar{z}')$

As above, T is necessarily singular, by Lemma 4. This rules out the last case above, as in the previous case, but also the first one as $E_0[Q''] \Downarrow_{\bar{z}}$.

To see that this is also absurd in the second case, observe that

$Q_i \Downarrow_{\bar{z}_i}$ and $z_i \in S'_I$. This entails that $T \Downarrow_{\bar{z}}$, hence a contradiction.

Second case: P' is not singular. By Lemma 4, we know that either P' is inactive, or $P' \Downarrow_{\bar{z}}$ for some z . As above, we rely on Theorem 2 to reason by cases according to the transition from P to P' . We note R_i for the process that we obtain, by induction, for each pair (P', Q_i) . We first show that for any $\tilde{z}_i, \tilde{x}, S, Q_i$ with $Q \stackrel{\hat{\mu}}{\Rightarrow} Q_i$, taking $\bar{R}_i$ accordingly with η' so

$$\bar{R}_i \stackrel{def}{=} \begin{cases} \bar{z}_i \mid z_i.R_i & \text{when } 0 \vdash_s P', Q_i \\ R + \sum_{y \in S} y(\tilde{y}'). \bar{z}_i & \text{when } 1 \vdash_s P', Q_i \end{cases}$$

and $F = (\nu \tilde{x})([\cdot] \mid \sum_{i \in I} \tau. \bar{R}_i)$, we have

$$F[P'] \stackrel{\cdot}{\approx}^1 F[Q_i] \quad (\star)$$

We have two cases, depending on the clause of the proposition that holds, by induction, for (P', Q_i) :

1. if clause 1) holds, then $F[Q_i] \rightarrow E_i[Q_i]$ with

$$E_i \stackrel{def}{=} (\nu \tilde{x})([\cdot] \mid \bar{R}_i).$$

This means there is A'_i with $F[P'] \Rightarrow A'_i$ and $A'_i \stackrel{\cdot}{\approx}^1 E_i[Q_i]$. As we do not have $E_i[Q_i] \Downarrow_{z_j}$ for $j \neq i$, we must have that

- either $A'_i \equiv E_i[P'']$ with $P' \Rightarrow P''$,
- or $A'_i \equiv (\nu \tilde{x})(P'' \mid \bar{z}_i)$ for some P'' .

However the second case is not possible as $E_i[Q_i] \rightarrow \Downarrow_{\bar{z}_i}$. Up to structural congruence, we can suppose that in $E_i[P'']$ and $E_i[Q_i]$, $\tilde{x}$ contain only names from $\text{fn}(P', Q_i)$. We are in a situation where we can apply the induction hypothesis, and this is in contradiction with $E_i[P''] \stackrel{\cdot}{\approx}^1 E_i[Q_i]$ as $P' \Rightarrow P''$.

2. if clause 2) holds, we have similarly that $F[P'] \rightarrow E_i[P']$. The same reasoning, with the role of P' and Q_i swapped, is sufficient to conclude.

We can now prove the desired result by looking at the possible cases for μ :

- when $\mu = \tau$ and $1 \vdash_s P'$, given fresh names $(z_i)_{i \in I}$, we define

$$R \stackrel{def}{=} \tau. \sum_{i \in I} \tau. (R_i + \sum_{y \in S'_I} y(\tilde{y}'). \bar{z}_i),$$

and $E_0 = (\nu \tilde{x})([\cdot] \mid R + \sum_{y \in S} y(\tilde{y}'). \bar{z})$. We reason by contradiction.

We have $E_0[P] \rightarrow F[P']$. Thus there should exist B' such that $B \Rightarrow B'$ and $F[P'] \stackrel{\cdot}{\approx}^1 B'$.

We have that P' is active and is not singular. This entails, by Lemma 4, that $P' \Downarrow_{\bar{z}'}$ for some $z' \in S'_I$ (P' is active because $1 \vdash_s P'$).

So $F[P'] \Downarrow_{\bar{z}_i}$ for all $i \in I$ but not $F[P'] \Downarrow_{\bar{z}}$. Therefore B' can be of two forms:

- $E_0[Q_i]$ for some $i \in I$ (with $Q' \Rightarrow Q_i$).
- $F[Q_i]$ for some $i \in I$ (with $Q' \Rightarrow Q_i$).

However, as $S'_I \subseteq S$, if $E_0[Q_i] \Downarrow_{\bar{z}_i}$, then we must have $E_0[Q_i] \Downarrow_{\bar{z}}$ which rules out the first case. Thus, $F[P'] \stackrel{\cdot}{\approx}^1 F[Q_i]$ which is absurd by $(\star)$.

- when $\mu = (\nu \tilde{y}^2) \bar{x}'(\tilde{y}^2, \tilde{y}^1)$, given fresh names $z', z'', (z_i)_{i \in I}$, we note $\tilde{y}^2 = y_1^2, \dots, y_n^2$ and $\tilde{y}^1 = y_1^1, \dots, y_m^1$. Then we set:

$$R \stackrel{def}{=} x'(\tilde{x}^2, \tilde{x}^1). (\bar{z}' \mid G \\ + [x_1^1 = y_1^1] \dots [x_m^1 = y_m^1] z'. (\sum_{i \in I} \tau. (\bar{z}_i \mid z_i.R_i)))$$

$$G \stackrel{def}{=} \sum_{j \leq n} \sum_{y'' \in \text{In}(P, Q)} [x_j^2 = y''] z'. \bar{z}''$$

and $E_0 = (\nu \tilde{x})([\cdot] R + \sum_{y \in S} y(\tilde{y}).\bar{z})$. Then $E_0[P] \rightarrow A$ with

$$A \equiv (\nu \tilde{x}, \tilde{x}^2)(P' | \bar{z}' | G + z'. (\sum_{i \in I} \tau. (\bar{z}_i | z_i. R_i)))$$

Thus there exists B such that $E_0[Q] \Rightarrow B$ and $A \stackrel{\cdot 1}{\approx} B$. We have $A \Downarrow_{\bar{z}}$, $A \Downarrow_{\bar{z}_i}$ for all $i \in I$ but not $A \Downarrow_{\bar{z}}$. Thus, B can only be of the form $(\nu \tilde{x}, \tilde{x}^2)(Q_i | G + z'. (\sum_{i \in I} \tau. (\bar{z}_i | z_i. R_i)))$ for some $i \in I$ (meaning that $Q' \stackrel{\mu}{\Rightarrow} Q_i$).

We have $A \rightarrow F[P']$ with $F[P'] \Downarrow_{\bar{z}_i}$ for all $i \in I$ but not $F[P'] \Downarrow_{\bar{z}}$. So we must have that $B \Rightarrow F[Q_{i'}]$ with $Q_i \Rightarrow Q_{i'}$ and $F[P'] \stackrel{\cdot 1}{\approx} F[Q_{i'}]$. This is absurd by $(\star)$.

The reasoning in the remaining cases is very similar to the proofs given above. We only sketch how the proof can be derived.

- when $\mu = x'(\tilde{y}')$, P is inactive, and P' is active. Given fresh names $(z_i)_{i \in I}$, we set

$$R \stackrel{def}{=} \bar{x}'(\tilde{y}') | \sum_{i \in I} \tau. (R_i + \sum_{y \in S'_i} y(\tilde{y}').\bar{z}_i)$$

and we can conclude.

- when $\mu = \tau$ and $0 \vdash_s P'$, given fresh names $(z_i)_{i \in I}$, we pose

$$R \stackrel{def}{=} \sum_{i \in I} \tau. (\bar{z}_i | z_i. R_i)$$

and we can conclude.

□

Appendix B. Completeness for WB-bisimulation

The proof of completeness for $\approx_{\text{wb}}^\sigma$ w.r.t. $\simeq^\sigma$ has the same overall structure as the proof for sequential bisimulation. Technically, however, because of the presence of continuation names, the details are different.

The following lemma is a direct consequence of the usage discipline of continuation names.

Lemma 15. *If $\sigma \vdash_{\text{wb}} P|Q$ then*

$$(\nu p)(\bar{p}(\tilde{a}) | p(\tilde{b}).P|Q) \approx_{\text{wb}}^\sigma P\{\tilde{a}/\tilde{b}\} | Q.$$

By soundness, this equivalence also holds for barbed equivalence at σ . Suppose $\sigma \vdash_{\text{wb}} P$. We recall the definition of $E_\sigma^{\tilde{x}_i}$, given in [Section 4.4.1](#):

With $\sigma = \xi, p_1 : \circ, q_1 : \mathbf{i}, \dots, p_{n-1} : \circ, q_{n-1} : \mathbf{i}, p_n : \circ$ and $\xi = \emptyset$ or $\xi = q : \mathbf{i}$, we define, for fresh q_n and $\tilde{x}_i$,

$$E_\sigma^{\tilde{x}_i} \stackrel{def}{=} (\nu \tilde{p}_i, \tilde{q}_i)([\cdot] \prod_{i \leq n} p_i(\tilde{y}).\bar{x}_i(\tilde{y}, q_i)).$$

When $[\sigma; P] \xrightarrow{(\nu \tilde{x})\bar{p}_1(\tilde{y})} [\sigma'; P']$, then with $\tilde{x}_i = x_i, \tilde{x}_i'$

$$q_n : \circ \vdash_{\text{wb}} E_\sigma^{\tilde{x}_i}[P] \xrightarrow{\tau} (\nu q_1, \tilde{x})(E_{\sigma'}^{\tilde{x}_i'}[P'] | \bar{x}_1(\tilde{y}, q_1))$$

By [Lemma 15](#), this internal transition does not change the behavior of the process. Thus, in the proof of [Proposition 9](#), we reason modulo this kind of reductions.

We can extend the definition of singular processes and derive results similar to [Lemmas 2, 3, 4](#):

Lemma 16. *Consider a clean stack $\sigma = p_0 : \circ, q_0 : \mathbf{i}, \dots, p_n : \circ$, and define*

$$\mathbf{0}_\sigma \stackrel{def}{=} (\nu x)\bar{x}(p_0) | \prod_{1 \leq i \leq n} q_{i-1}(\tilde{y}).\bar{p}_i(\tilde{y}).$$

Then for any singular process P with $\sigma \vdash_{\text{wb}} P$, we have $P \approx_{\text{wb}}^\sigma \mathbf{0}_\sigma$. Similarly, if P does not contain any input-controlled name, then P is singular iff $P \stackrel{\cdot \sigma}{\approx} \mathbf{0}_\sigma$.

Under the same hypotheses, this lemma implies that in absence of input-controlled names, P is not singular iff $E_\sigma^{\tilde{x}_i}[P] \Downarrow_{\bar{z}}$ for some z .

The following proposition is the counterpart, for WB-bisimulation, of [Proposition 5](#). It refers to the stratification of WB-bisimulation, which is defined like the stratification of sequential bisimulation.

Proposition 9. For all $n \geq 0$, $\sigma = \xi, \sigma_0$ with $\xi = \emptyset$ or $\xi = p : i$, if $\sigma \vdash_{\text{wb}} P, Q$ and $P \not\approx_{\text{wb}}^{\sigma, n} Q$, then for any fresh names $\tilde{x}_j$ there exists R such that for any fresh z and any $\tilde{x}, S$ with $\tilde{x} \subseteq \text{fn}_0(P, Q)$ and $\text{fn}_0(P, Q) \cup \tilde{x}_j \subseteq S$, we define:

$$\begin{aligned} \bar{R} &\stackrel{\text{def}}{=} \begin{cases} \bar{z}(p) | z(p).R & \text{when } \xi = p : i \\ R + \sum_{y \in S} y(\tilde{y}', q). \bar{z}(q) & \text{when } \xi = \emptyset \end{cases} \\ E_0 &\stackrel{\text{def}}{=} \begin{cases} (\nu \tilde{x}, p)(E_{\sigma}^{\tilde{x}_j} | \bar{R}) & \text{when } \xi = p : i \\ (\nu \tilde{x})(E_{\sigma}^{\tilde{x}_j} | \bar{R}) & \text{when } \xi = \emptyset \end{cases} \end{aligned}$$

and we have:

1. either $E_0[P'] \not\approx^{q_n, \circ} E[Q]$ for all P' such that $P \Rightarrow P'$,
2. or $E_0[P] \not\approx^{q_n, \circ} E[Q']$ for all Q' such that $Q \Rightarrow Q'$.

The proof of this proposition is structured like the proof of Proposition 5. We therefore do not provide all details, but instead highlight the technical points that are specific.

Proof. We reason by induction on n . For $n = 0$, there is nothing to prove. For $n > 0$, suppose that $P \not\approx_{\text{wb}}^{\sigma, n} Q$. We may therefore suppose that there exists μ such that $[\sigma; P] \xrightarrow{\mu} [\sigma'; P']$ and for all Q' with $Q \xRightarrow{\hat{\mu}} Q'$ (or when $\mu = x(\tilde{y}, p)$, all Q' with

$Q | (\nu q)(\bar{x}(\tilde{y}, q) | q(\tilde{y})). \bar{p}(\tilde{y}')) \Rightarrow Q'$ for any fresh q), we have $P' \not\approx_{\text{wb}}^{\sigma', n-1} Q'$. We note $\{Q_i\}$ for $i \in I$ the set of all such Q' . This set is finite by hypothesis. We also write

$$S'_I = \bigcup_{i \in I} \text{fn}_0(Q_i) \cup \text{fn}_0(P') \cup \tilde{x}_j, \text{ which is also a finite set.}$$

We show that clause 1) holds. (in the symmetrical case, where we consider a transition from $[\eta; Q]$, we establish the second clause).

We distinguish two cases, according to whether P' is singular or not. The first case is proved with the same reasoning as for sequentiality. For the second case, we have to be more cautious as the presence of $E_{\sigma}^{\tilde{x}_j}$ enables an additional reduction. However, this reduction does not affect the behavior, thanks to Lemma 15.

First case: P' is singular. Since P' is singular, $\xi' = \emptyset$. For all i , as $P' \not\approx_{\text{wb}} \sigma' Q_i$, we have $E_{\sigma'}^{\tilde{x}_j}[Q_i] \Downarrow_{\tilde{z}_i}$ for some $z_i \in S'_I$.

By Theorem 5, there are three possible cases:

- when $\mu = \tau$ and $\sigma' = \sigma$. We set $R \stackrel{\text{def}}{=} 0$. We then have to show that for any $z, \tilde{x}, S, Q'$ with $Q \Rightarrow Q'$, taking $\bar{R} = R + \sum_{y \in S} y(\tilde{y}', q). \bar{z}(q)$ (since $\xi = \emptyset$) and $E_0 = (\nu \tilde{x})(E_{\sigma}^{\tilde{x}_j} | \bar{R})$, we have $E_0[P] \not\approx^{q_n, \circ} E_0[Q']$. We reason by contradiction.

We observe $E_0[P] \rightarrow E_0[P']$ which is singular. Thus there should exist T such that $E_0[Q'] \Rightarrow T$ and $E_0[P'] \approx^{q_n, \circ} T$, which entails that T is singular.

We observe then that T can be of 3 forms:

- $E_0[Q_i]$ for some $i \in I$ (with $Q' \Rightarrow Q_i$)
- $(\nu \tilde{x}, q_1)(Q'' | \bar{x}_1(\tilde{y}, q_1) | \bar{R})$
- $(\nu \tilde{x}, q)(Q'' | \bar{z}(q))$

We have one more case than in the proof of completeness for sequential bisimilarity due to the distinction between continuation names and the other output-controlled names. Since T is singular, the latter two are impossible. For the first case, we have that $S'_I \subseteq S$ and $E_{\sigma'}^{\tilde{x}_j}[Q_i] \Downarrow_{\tilde{z}_i}$ for some $z_i \in S'_I$, so $T \Downarrow_{\tilde{z}}$, which contradicts the fact that T is singular.

- when $\mu = p(\tilde{y}')$, given a fresh name z' , we define

$$R \stackrel{\text{def}}{=} \bar{p}(\tilde{y}') | \sum_{y \in S'_I} y(\tilde{y}', q). \bar{z}'(q).$$

The proof is similar to the one for the completeness for sequential bisimilarity, with, as above, one more case for T that is handled similarly.

- when $\mu = x'(\tilde{y}', p_0)$

Given fresh names z' and x_0 , we define

$$\begin{aligned} R &\stackrel{\text{def}}{=} (\nu p_0)(\bar{x}'(\tilde{y}', p_0) | p_0(\tilde{y}'). \bar{x}_0(\tilde{y}', q_0)) \\ &\quad | \sum_{y \in S'_I} y(\tilde{y}', q). \bar{z}'(q) \end{aligned}$$

We then have to show that for any $z, \tilde{x}, S, Q'$ with $Q \Rightarrow Q'$, taking $\bar{R} = z(q_0) | \bar{z}(q_0)$ and $E_0 = (\nu \tilde{x}, q_0)(E_{\sigma}^{\tilde{x}_j} | \bar{R})$, we have

$$E_0[P] \not\approx^{q_n, \circ} E_0[Q'].$$

We reason by contradiction and note $E_1 \stackrel{\text{def}}{=} (\nu \tilde{x})(E_{\sigma'}^{\tilde{x}_j} | \sum_{y \in S'_I} y(\tilde{y}', q))$. First, we can see that $E_0[Q'] \approx_s^{q_n, \circ} (\nu \tilde{x}, q_0)(E_{\sigma}^{\tilde{x}_j}[Q'] | z(q_0). R' | \bar{z}(q_0))$ where R' is the process obtained by applying the law of Lemma 7 to $\bar{x}'(\tilde{y}', p_0)$ in R .

We have $E_0[P] \rightarrow E_1[P']$ which is singular. Thus there should exist some T such that

$(\nu \tilde{x}, q_0)(E_{\sigma}^{\tilde{x}j}[Q']|z(q_0).R'|\bar{z}\langle q_0 \rangle) \Rightarrow T$ and $E_1[P'] \approx^{q_n:\circ} T$. We observe that T can be of 4 forms:

- $(\nu \tilde{x}, q_0)(E_{\sigma}^{\tilde{x}j}[Q'']|z(q_0).R'|\bar{z}\langle q_0 \rangle)$ (with $Q' \Rightarrow Q''$)
- $E_1[Q_i]$ for some $i \in I$ (with $Q'|(vr_0)(\bar{x}'\langle \tilde{y}', r_0 \rangle|r_0(\tilde{y}').\bar{p}_0(\tilde{y}')) \Rightarrow Q_i$)
- $(\nu \tilde{x}, q_1)(Q''|\bar{x}_1\langle \tilde{y}, q_1 \rangle)$
- $(\nu \tilde{x}, q)(Q''|\bar{z}'\langle q \rangle)$

As above, T is necessarily singular. This immediately rules out all but the second case. For the latter case, $E_{\sigma'}^{\tilde{x}_0, \tilde{x}j}[Q_i] \Downarrow_{\bar{z}_i}$ for some $z_i \in S'_I$, so we have that $T \Downarrow_{\bar{z}_i}$, hence a contradiction.

Second case: if P' is not singular. We note R_i for the process we obtain, by induction, for each pair (P', Q_i) . As for the proof of completeness for sequential bisimilarity, we show that for any $\tilde{z}_i, \tilde{x}, S, Q_i$ with $Q \xRightarrow{\hat{\mu}} Q_i$, we define

$$\begin{aligned} \bar{R}_i &\stackrel{def}{=} \begin{cases} \bar{z}_i\langle p \rangle | z_i(p).R_i & \text{when } \xi' = p : i \\ R_i + \sum_{y \in S} y(\tilde{y}', q).\bar{z}_i\langle q \rangle & \text{when } \xi' = \emptyset \end{cases} \\ F &\stackrel{def}{=} \begin{cases} (\nu \tilde{x}, p)(E_{\sigma}^{\tilde{x}j} | \sum_{i \in I} \tau.\bar{R}_i) & \text{when } \xi' = p : i \\ (\nu \tilde{x})(E_{\sigma}^{\tilde{x}j} | \sum_{i \in I} \tau.\bar{R}_i) & \text{when } \xi' = \emptyset \end{cases} \end{aligned}$$

and we have

$$F[P'] \approx^{q_n:\circ} F[Q_i]. \quad (\star)$$

We have two cases depending on the clause of the proposition that holds, by induction, for (P', Q_i) :

1. if clause 1) holds, then $F[Q_i] \rightarrow E_i[Q_i]$ with

$$E_i \stackrel{def}{=} (\nu \tilde{x})(E_{\sigma}^{\tilde{x}j} | \bar{R}_i)$$

This means there is A'_i with $F[P'] \Rightarrow A'_i$ and $A'_i \approx^{q_n:\circ} E_i[Q_i]$. As we do not have $E_i[Q_i] \Downarrow_{z_j}$ for $j \neq i$, we must have that

- either $A'_i \equiv E_i[P'']$ with $P' \Rightarrow P''$,
- $A'_i \equiv (\nu \tilde{x}, \tilde{x}', q_1)(\tilde{P}|\bar{x}_1\langle \tilde{y}', q_1 \rangle | \bar{R}_i)$ with $P' \Rightarrow P'' \xrightarrow{(\nu \tilde{x}')\bar{p}_1\langle \tilde{y}' \rangle} \tilde{P}$
- or $A'_i \equiv (\nu \tilde{x}, q)(P''|\bar{z}_i\langle q \rangle)$ for some P'' .

However the last case is not possible as $E_i[Q_i] \not\Downarrow_{\bar{z}_i}$. Because of [Lemma 15](#), the second case is barbed bisimilar to the first one. Up to structural congruence, we can suppose that in $E_i[P'']$ and $E_i[Q_i]$, $\tilde{x}$ contain only names from $\text{fn}(P', Q_i)$. We are in a situation where we can apply the induction hypothesis, and this is in contradiction with $E_i[P''] \approx^{q_n:\circ} E_i[Q_i]$ as $P' \Rightarrow P''$.

2. if clause (2) holds, we have similarly that $F[P'] \rightarrow E_i[P']$. The same reasoning, with the role of P' and Q_i swapped, is sufficient to conclude.

We can now examine the different possibilities for the transition along μ .

- $\mu = x'\langle \tilde{y}', p_0 \rangle$, then $\xi = q_0 : i$ and $\sigma' = p_0 : o, \sigma$. Given fresh names $(z_i)_{i \in I}$, we define

$$\begin{aligned} R &\stackrel{def}{=} (\nu p_0)(\bar{x}'\langle \tilde{y}', p_0 \rangle | p_0(\tilde{y}').\bar{x}_0\langle \tilde{y}', q_0 \rangle) \\ &\quad | \sum_{i \in I} \tau.(R_i + \sum_{y \in S'_I} y(\tilde{y}', q).\bar{z}_i\langle q \rangle) \end{aligned}$$

and call R' the process obtained by applying the law of [Lemma 7](#) to $\bar{x}_0\langle \tilde{y}', q_0 \rangle$ in R . (Thus $R \cong^{q_0:\circ} R'$).

We then have to show that for any $z, \tilde{x}, S, Q'$ with $Q \Rightarrow Q'$, we have $E_0[P] \approx^{q_n:\circ} B$ where

$$B \stackrel{def}{=} (\nu \tilde{x}, q_0)(E_{\sigma}^{\tilde{x}j}[Q']|z(q_0).R'|\bar{z}\langle q_0 \rangle).$$

We reason by contradiction. We have $E_0[P] \rightarrow F[P']$ with $F[P'] \not\Downarrow_{\bar{z}}$.

Thus, there exists T such that $B \Rightarrow T$ and $F[P'] \approx^{q_n:\circ} T$. As P' is not singular, $E_{\sigma'}^{\tilde{x}_0, \tilde{x}j}[P'] \Downarrow_{\bar{z}'}$ for some $z' \in S'_I$, so $F[P'] \Downarrow_{\bar{z}_i}$ for all z_i . This implies that $T \equiv F[Q_i]$ with $Q'|(vr_0)(\bar{x}'\langle \tilde{y}', r_0 \rangle|r_0(\tilde{y}').\bar{p}_0(\tilde{y}')) \Rightarrow Q_i$.

This is in contradiction with $(\star)$.

- $\mu = (\nu \tilde{y}^2, q)\bar{x}'\langle \tilde{y}', q \rangle$. We note $\tilde{y}^2 = y_1^2, \dots, y_n^2$ and $\tilde{y}' = \tilde{y}^2, y_1^1, \dots, y_m^1$. Given fresh names $z', z'', (z_i)_{i \in I}$, we define

$$\begin{aligned} R &\stackrel{def}{=} x'(\tilde{x}^2, \tilde{x}^1, q_1).(\bar{z}'\langle q_1 \rangle | R_1) \\ R_1 &\stackrel{def}{=} \sum_{j \leq n} \sum_{y'' \in \text{fn}_0(P, Q)} [x_j^2 = y''] z'(q).\bar{z}''\langle q \rangle \end{aligned}$$

$$+ [x_1^1 = y_1^1] \dots [x_m^1 = y_m^1] z'(q). \sum_{i \in I} \tau. (\bar{z}_i \langle q \rangle | z_i(q'). R_i)$$

The reasoning is similar to the proof of [Proposition 5](#).

- $\mu = (v\tilde{y}^2)\bar{p}_1 \langle \tilde{y}' \rangle$. We note $\tilde{y}^2 = y_1^2, \dots, y_n^2$ and $\tilde{y}' = \tilde{y}^2, y_1^1, \dots, y_m^1$. Given fresh names $z', z'', (z_i)_{i \in I}$, we define

$$\begin{aligned} R &\stackrel{def}{=} x_1(\tilde{x}^2, \tilde{x}^1, q_1).(\bar{z}' \langle q_1 \rangle | R_1) \quad \text{with} \\ R_1 &\stackrel{def}{=} \sum_{j \leq n} \sum_{y'' \in \text{fn}_0(P, Q)} [x_j^2 = y''] z'(q). \bar{z}'' \langle q \rangle \\ &\quad + [x_1^1 = y_1^1] \dots [x_m^1 = y_m^1] z'(q). \sum_{i \in I} \tau. (\bar{z}_i \langle q \rangle | z_i(q'). R_i) \end{aligned}$$

Here we make use of $E_{\sigma}^{\tilde{x}^j}$ to test the output at p_1 by testing the output at x_1 . We conclude as before.

- $\mu = \tau$ and $\xi = \emptyset$. Given fresh names $(z_i)_{i \in I}$, we define $R \stackrel{def}{=} \sum_{i \in I} \tau. (R_i + \sum_{y \in S'_i} y(\tilde{y}', q). \bar{z}_i \langle q \rangle)$, and we can conclude.
- $\mu = \tau$ and $\xi = p : i$. Given fresh names $(z_i)_{i \in I}$, we define $R \stackrel{def}{=} \sum_{i \in I} \tau. (\bar{z}_i \langle p \rangle | z_i(p). R_i)$, and we can conclude.
- $\mu = p \langle \tilde{y}' \rangle$. Given fresh names $(z_i)_{i \in I}$, we define

$$R \stackrel{def}{=} \bar{p} \langle \tilde{y}' \rangle | \sum_{i \in I} \tau. (R_i + \sum_{y \in S'_i} y(\tilde{y}', q). \bar{z}_i \langle q \rangle)$$

and we can conclude.

□

Appendix C. Examples with wb-bisimulation ([Section 4.4.2](#))

C.1. A simplified example

We discuss below a simplified example, which exposes the main difficulties that arise when studying the well-bracketed state change example. The primary simplification consists in using single-use functions, i.e., without replication. As a consequence, the two calls to the external function are split into two separate functions. In ML, this corresponds to the terms N and L below:

$$\begin{array}{ll} N &\stackrel{def}{=} \text{let } x = \text{ref } 0 \text{ in } (N_1, N_2) \\ N_1 &\stackrel{def}{=} \text{fun } f \rightarrow x := 1; f (); !x \\ N_2 &\stackrel{def}{=} \text{fun } f \rightarrow x := 0; f (); x := 1 \\ \\ L &\stackrel{def}{=} (L_1, L_2) \\ L_1 &\stackrel{def}{=} \text{fun } f \rightarrow f (); 1 \\ L_2 &\stackrel{def}{=} \text{fun } f \rightarrow f (); () \end{array}$$

with the constraint that each term may only be called once. (Such a simplification would not work on the well-bracketed state change example, as M_1 and M_2 become equivalent even dropping well-bracketing if used at most once.) Intuitively,

$$\text{let}(h_1, h_2) = N \text{ in fun } f \rightarrow h_2 f; h_1 f$$

is similar to a single use of M_1 , and the same for L and M_2 .

The $\Lambda\pi$ translation of the above terms is as follows:

$$\begin{aligned} \llbracket N \rrbracket_{p_0} &\stackrel{def}{=} (v\ell, x, y)(\bar{\ell} \langle 0 \rangle | \bar{p}_0 \langle x, y \rangle | P_1 | P_2) \\ P_1 &\stackrel{def}{=} x(z, p). \text{wx}_{\ell} \langle 1 \rangle. (vp')(\bar{z} \langle \star, p' \rangle | p'. \text{xe}_{\ell}(n). \bar{p} \langle n \rangle) \\ P_2 &\stackrel{def}{=} y(z, q). \text{wx}_{\ell} \langle 0 \rangle. (vq')(\bar{z} \langle \star, q' \rangle | q'. \text{wx}_{\ell} \langle 1 \rangle. \bar{q}) \\ \llbracket L \rrbracket_{p_0} &\stackrel{def}{=} (vx, y)(\bar{p}_0 \langle x, y \rangle | Q_1 | Q_2) \\ Q_1 &\stackrel{def}{=} x(z, p). (vp')(\bar{z} \langle \star, p' \rangle | p'. \bar{p} \langle 1 \rangle) \\ Q_2 &\stackrel{def}{=} y(z, q). (vq')(\bar{z} \langle \star, q' \rangle | q'. \bar{q}) \end{aligned}$$

Indeed, the translation of an ML program is parametrised upon a continuation name (here p_0).

We introduce a slight abuse of notation, and write $\llbracket N \rrbracket_{p_0} \xrightarrow{(vx, y)\bar{p}_0 \langle x, y \rangle} \llbracket N \rrbracket$, and similarly for notation $\llbracket L \rrbracket$. Below we show that $\llbracket N \rrbracket$ and $\llbracket L \rrbracket$ are equivalent. First however we note that well-bracketing is necessary for this. Sequentiality alone is not sufficient: under the type system for sequentiality we can observe the following trace from $\llbracket N \rrbracket$:

$$\llbracket N \rrbracket \xrightarrow{x(\star, p)} \xrightarrow{(vp')\bar{z} \langle \star, p' \rangle} \xrightarrow{y(\star, q)} \xrightarrow{\bar{z} \langle \star, q' \rangle} \xrightarrow{p'} \bar{p} \langle 0 \rangle$$

This trace is not well-bracketed: function x is called first, but the continuation p is returned before q . Process $\llbracket L \rrbracket$ may not produce such a trace — it may only emit 1. However, the above non-well-bracketed trace is the only trace from $\llbracket N \rrbracket$ that may produce 0.

We prove $\llbracket N \rrbracket \approx_{\text{wb}}^{\emptyset} \llbracket L \rrbracket$ by defining a relation, and relying on up-to techniques for WB-bisimulation. For that, we use the abbreviations

$$\begin{array}{ll} P'_1 \stackrel{\text{def}}{=} p'.\text{re}_c(n).\bar{p}\langle n \rangle, & Q'_1 \stackrel{\text{def}}{=} p'.\bar{p}\langle 1 \rangle, \\ P'_2 \stackrel{\text{def}}{=} q'.\text{wr}_c\langle 1 \rangle.\bar{q} & Q'_2 \stackrel{\text{def}}{=} q'.\bar{q} \end{array}$$

This allows us to define the following relation, called $\mathcal{R}$:

$$\begin{array}{l} \{(\emptyset, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1 | P_2), Q_1 | Q_2), \\ ((p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P'_1 | P_2), Q'_1 | Q_2), \\ ((q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1 | P'_2), Q_1 | Q'_2), \\ ((q' : i, q : o, p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P'_1 | P'_2), Q'_1 | Q'_2), \\ (\emptyset, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_2), Q_2), \\ (\emptyset, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1), Q_1), \\ ((p' : i, p : o, q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P'_1 | P'_2), Q'_1 | Q'_2), \\ ((p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P'_1), Q'_1), \\ ((q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P'_2), Q'_2), \\ ((q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P'_2), Q'_2), \\ (\emptyset, (\nu \ell)\bar{\ell}\langle 1 \rangle, \mathbf{0}) \} \end{array}$$

$\mathcal{R}$ is a wb-bisimulation up-to deterministic reduction and up-to static context. Ensuring well-bracketing restricts transitions for the two triplets below:

$$((q' : i, q : o, p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P'_1 | P'_2), Q'_1 | Q'_2)$$

$$((p' : i, p : o, q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P'_1 | P'_2), Q'_1 | Q'_2)$$

Indeed, we may only test the input on q' (resp. p') while both are unguarded in both cases. This is crucial for the first triplet as allowing the input on p' on the first process would lead to a process (after a deterministic reduction) that can send $\bar{p}0$, while the second one cannot emit 0.

Without up-to techniques. We define some additional relations, which allow us to show the improvement brought by up-to techniques. For this, we introduce the following notations for processes related to P_1 and Q_1 respectively:

$$\begin{array}{llll} P_1^a \stackrel{\text{def}}{=} \text{wr}_c\langle 1 \rangle.P_1^b & P_1^b \stackrel{\text{def}}{=} (\nu p')(\bar{z}\langle \star, p' \rangle | P_1^c) & P_1^c \stackrel{\text{def}}{=} p'.P_1^d & P_1^d \stackrel{\text{def}}{=} \text{re}_c(n).\bar{p}\langle n \rangle \\ Q_1^c \stackrel{\text{def}}{=} p'.\bar{p}\langle 1 \rangle & Q_1^b \stackrel{\text{def}}{=} (\nu p')(\bar{z}\langle \star, p' \rangle | Q_1^c) & & \end{array}$$

and similarly for processes P_2 and Q_2 . We can remark that P_i^c is the same as P_i' introduced above to define $\mathcal{R}$.

$$\begin{array}{l} \mathcal{R}' \stackrel{\text{def}}{=} \{ (p : o, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1^a | P_2), Q_1^a | Q_2), \\ (p : o, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1 | P_2^a), Q_1 | Q_2^b), \\ ((q : o, p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^c | P_2^a), Q_1^c | Q_2^b), \\ (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^d | P_2), \bar{p}\langle 1 \rangle | Q_2), \\ (q : o, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1 | P_2^d), Q_1 | \bar{q}), \\ ((p : o, q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1^a | P_2^c), Q_1^b | Q_2^c), \\ ((q : o, p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1^c | P_2^d), Q_1^c | \bar{q}), \\ (q : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_2^a), Q_2^b), \\ (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^a), Q_1^b), \\ ((p : o, q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^d | P_2^c), \bar{p}\langle 1 \rangle | Q_2^c), \\ (q : o, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_2^d), \bar{q}), \\ (q : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_2^d), \bar{q}), \\ (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^d), \bar{p}\langle 1 \rangle) \} \end{array}$$

$$\begin{aligned}
\mathcal{R}'' \stackrel{def}{=} \{ & (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^b | P_2), Q_1^b | Q_2), \\
& (p : o, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1 | P_2^b), Q_1 | Q_2^b), \\
& ((q : o, p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_1^c | P_2^b), Q_1^c | Q_2^b), \\
& (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | \bar{p}\langle 1 \rangle | P_2), \bar{p}\langle 1 \rangle | Q_2), \\
& (q : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1 | \bar{q}), Q_1 | \bar{q}), \\
& ((p : o, q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^b | P_2^c), Q_1^b | Q_2^c), \\
& ((q : o, p' : i, p : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^c | \bar{q}), Q_1^c | \bar{q}), \\
& (q : o, (\nu \ell)(\bar{\ell}\langle 0 \rangle | P_2^b), Q_2^b), \\
& (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | P_1^b), Q_1^b), \\
& ((p : o, q' : i, q : o), (\nu \ell)(\bar{\ell}\langle 1 \rangle | \bar{p}\langle 1 \rangle | P_2^c), \bar{p}\langle 1 \rangle | Q_2^c), \\
& (q : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | \bar{q}), \bar{q}), \\
& (p : o, (\nu \ell)(\bar{\ell}\langle 1 \rangle | \bar{p}\langle 1 \rangle), \bar{p}\langle 1 \rangle) \\
& \}
\end{aligned}$$

We see that $\mathcal{R}$, $\mathcal{R}'$ and $\mathcal{R}''$ are comparable in size. We have:

- $\mathcal{R} \uplus \mathcal{R}'$ is a wb-bisimulation up-to static context.
- $\mathcal{R} \uplus \mathcal{R}''$ is a wb-bisimulation up-to deterministic τ .
- $\mathcal{R} \uplus \mathcal{R}' \uplus \mathcal{R}''$ is a wb-bisimulation.

C.2. Well-bracketed state change

We present the proof for the example described in [Section 4.4.2](#). We need to introduce some notations in order to reason about the sub-processes which are generated by calls to the functions in the encodings of M_1 and M_2 . To simplify notation, we assume that names extruded from (νq) (resp. (νr)) are in the set $\{q_i \mid i \in I\}$ (resp. $\{r_j \mid j \in J\}$) for some set I (resp. J). Continuation names received on x will be noted p_i or p_j depending on which prefix they are underneath.

$$\begin{aligned}
P_i^1 &= q_i.(\nu r)(\bar{y}\langle r \rangle | r.\bar{p}_i\langle 1 \rangle) & P_j^2 &= r_j.\bar{p}_j\langle 1 \rangle & Q_0 &= !x(y, p).\text{wr}_{\ell}\langle 0 \rangle.(\nu q)(\bar{y}\langle q \rangle | q.\text{wr}_{\ell}\langle 1 \rangle.(\nu r)(\bar{y}\langle r \rangle | r.\text{re}_{\ell}(n).\bar{p}\langle n \rangle)) \\
Q_i^1 &= q_i.\text{wr}_{\ell}\langle 1 \rangle.(\nu r)(\bar{y}\langle r \rangle | r.\text{re}_{\ell}(n).\bar{p}_i\langle n \rangle) & Q_j^2 &= r_j.\text{re}_{\ell}(n).\bar{p}_j\langle n \rangle
\end{aligned}$$

Given an ordered list s of the indices in $I \uplus J$, we write $\sigma(s)$ for the stack defined inductively as follows:

$$\begin{aligned}
\sigma(\emptyset) &= \emptyset & \sigma(i; s) &= q_i : i, p_i : o, \sigma(s) & \sigma(j; s) &= r_j : i, p_j : o, \sigma(s)
\end{aligned}$$

Using the above notations, we can then define the relation:

$$\begin{aligned}
\mathcal{R} \stackrel{def}{=} \{ & (\emptyset, P, Q), \quad (\emptyset, P, (\nu \ell)(\bar{\ell}\langle 1 \rangle | Q_0)), \\
& (\sigma(i; s), P | \prod_{i \in I} P_i^1 | \prod_{j \in J} P_j^2, \\
& (\nu \ell)(\bar{\ell}\langle 0 \rangle | Q_0 | \prod_{i \in I} Q_i^1 | \prod_{j \in J} Q_j^2)), \\
& (\sigma(i; s), P | \prod_{i \in I} P_i^1 | \prod_{j \in J} P_j^2, \\
& (\nu \ell)(\bar{\ell}\langle 1 \rangle | Q_0 | \prod_{i \in I} Q_i^1 | \prod_{j \in J} Q_j^2)), \\
& (\sigma(j; s), P | \prod_{i \in I} P_i^1 | \prod_{j \in J} P_j^2, \\
& (\nu \ell)(\bar{\ell}\langle 1 \rangle | Q_0 | \prod_{i \in I} Q_i^1 | \prod_{j \in J} Q_j^2)) \\
& \}
\end{aligned}$$

$\mathcal{R}$ is a wb-bisimulation up-to deterministic τ and up-to static context.

References

- [1] D. Sangiorgi, D. Walker, *The Pi-Calculus - a Theory of Mobile Processes*, Cambridge University Press, 2001.
- [2] B.C. Pierce, D. Sangiorgi, Typing and subtyping for mobile processes, *Math. Struct. Comput. Sci.* 6 (5) (1996) 409–453.
- [3] N. Kobayashi, B.C. Pierce, D.N. Turner, Linear and the pi-calculus, *ACM Trans. Program. Lang. Syst.* 21 (5) (1999) 914–947. <https://doi.org/10.1145/330249.330251>
- [4] K. Honda, V.T. Vasconcelos, M. Kubo, Language primitives and type discipline for structured communication-based programming, in: C. Hankin (Ed.), *ESOP'98, 7th European Symposium on Programming, 1998, Proceedings, Vol. 1381 of Lecture Notes in Computer Science*, Springer, 1998, pp. 122–138. <https://doi.org/10.1007/BFb0053567>
- [5] D. Ancona, V. Bono, M. Bravetti, J. Campos, G. Castagna, P. Deniéou, S.J. Gay, N. Gesbert, E. Giachino, R. Hu, E.B. Johnsen, F. Martins, V. Mascardi, F. Montesi, R. Neykova, N. Ng, L. Padovani, V.T. Vasconcelos, N. Yoshida, Behavioral types in programming languages, *Found. Trends Program. Lang.* 3 (2-3) (2016) 95–230. <https://doi.org/10.1561/25000000031>
- [6] S. Gay, V. Vasconcelos, *Session Types*, Cambridge University Press, 2025.
- [7] N. Kobayashi, A partially deadlock-free typed process calculus, *ACM Trans. Program. Lang. Syst.* 20 (2) (1998) 436–482. <https://doi.org/10.1145/276393.278524>
- [8] N. Kobayashi, A type system for lock-free processes, *Inf. Comput.* 177 (2) (2002) 122–159. <https://doi.org/10.1006/inco.2002.3171>
- [9] Y. Deng, D. Sangiorgi, Ensuring termination by typability, *Inf. Comput.* 204 (7) (2006) 1045–1082. <https://doi.org/10.1016/j.ic.2006.03.002>
- [10] N. Yoshida, M. Berger, K. Honda, Strong normalisation in the Pi-calculus, *Inf. Comput.* 191 (2) (2004) 145–202. <https://doi.org/10.1016/j.ic.2003.08.004>
- [11] N. Kobayashi, D. Sangiorgi, A hybrid type system for lock-freedom of mobile processes, *ACM Trans. Program. Lang. Syst.* 32 (5) (2010) 16:1–16:49. <https://doi.org/10.1145/1745312.1745313>
- [12] K. Honda, N. Yoshida, On reduction-based process semantics, *TCS* 152 (2) (1995) 437–486.
- [13] D. Sangiorgi, D. Walker, Some results on barbed equivalences in Pi-calculus, in: *Proc. CONCUR '01, Vol. 2154 of Lecture Notes in Computer Science*, Springer Verlag, 2001.
- [14] R. Milner, Functions as processes, *Math. Struct. Comput. Sci.* 2 (2) (1992) 119–141. <https://doi.org/10.1017/S0960129500001407>
- [15] D. Sangiorgi, An investigation into functions as processes, in: *Proc. Ninth International Conference on the Mathematical Foundations of Programming Semantics (MFPS'93)*, Vol. 802 of *Lecture Notes in Computer Science*, Springer, 1993, pp. 143–159.
- [16] D. Sangiorgi, The name discipline of uniform receptiveness, *Theor. Comput. Sci.* 221 (1-2) (1999) 457–493.
- [17] D. Hirschhoff, E. Prebet, D. Sangiorgi, On the representation of references in the pi-calculus, in: I. Konnov, L. Kovács (Eds.), *31st International Conference on Concurrency Theory, CONCUR 2020, Vol. 171 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik*, 2020, pp. 34:1–34:20. <https://doi.org/10.4230/LIPIcs.CONCUR.2020.34>
- [18] R.M. Amadio, I. Castellani, D. Sangiorgi, On bisimulations for the asynchronous pi-calculus, *Theor. Comput. Sci.* 195 (2) (1998) 291–324. [https://doi.org/10.1016/S0304-3975\(97\)00223-5](https://doi.org/10.1016/S0304-3975(97)00223-5)
- [19] R. Milner, *The Polyadic π -Calculus: A Tutorial*, Technical Report ECS-LFCS-91–180, LFCS, 1991. Also in *Logic and Algebra of Specification*, ed. F.L. Bauer, W. Brauer and H. Schwichtenberg, Springer Verlag, 1993.
- [20] M. Merro, D. Sangiorgi, On asynchrony in name-passing calculi, *Math. Struct. Comp. Sci.* 14 (5) (2004) 715–767.
- [21] J.M.E. Hyland, C.H.L. Ong, On full abstraction for PCF: I, II, and III, *Inf. Comput.* 163 (2) (2000) 285–408. <https://doi.org/10.1006/inco.2000.2917>
- [22] A. Ahmed, D. Dreyer, A. Rossberg, State-dependent representation independence, in: Z. Shao, B.C. Pierce (Eds.), *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, ACM, 2009*, pp. 340–353. <https://doi.org/10.1145/1480881.1480925>
- [23] D. Dreyer, G. Neis, L. Birkedal, The impact of higher-order state and control effects on local relational reasoning, *J. Funct. Program.* 22 (4-5) (2012) 477–528. <https://doi.org/10.1017/S095679681200024X>
- [24] G. Jaber, SyTeCi: automating contextual equivalence for higher-order programs with references, *Proc. ACM Program. Lang.* 4 (POPL) (2020) 59:1–59:28. <https://doi.org/10.1145/3371127>
- [25] F. Levi, D. Sangiorgi, Controlling interference in ambients, in: M.N. Wegman, T.W. Reps (Eds.), *POPL 2000, Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, 2000, ACM, 2000*, pp. 352–364. <https://doi.org/10.1145/325694.325741>
- [26] M. Berger, K. Honda, N. Yoshida, Sequentiality and the pi-Calculus, in: S. Abramsky (Ed.), *Typed Lambda Calculi and Applications, 5th International Conference, TLCA 2001, Proceedings, Vol. 2044 of Lecture Notes in Computer Science*, Springer, 2001, pp. 29–45. https://doi.org/10.1007/3-540-45413-6_7
- [27] M. Hennessy, J. Rathke, Typed behavioural equivalences for processes in the presence of subtyping, *Math. Struct. Comput. Sci.* 14 (5) (2004) 651–684. <https://doi.org/10.1017/S0960129504004281>
- [28] A. Igarashi, N. Kobayashi, Type reconstruction for linear-calculus with I/O subtyping, *Inf. Comput.* 161 (1) (2000) 1–44. <https://doi.org/10.1006/inco.2000.2872>
- [29] J. Laird, Full abstraction for functional languages with control, in: *Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science, 1997, IEEE Computer Society, 1997*, pp. 58–67. <https://doi.org/10.1109/LICS.1997.614931>
- [30] K. Honda, Processes and games, *Electron. Notes Theor. Comput. Sci.* 71 (2002) 40–69. [https://doi.org/10.1016/S1571-0661\(05\)82528-9](https://doi.org/10.1016/S1571-0661(05)82528-9)
- [31] L. Skorstengaard, D. Devriese, L. Birkedal, StkTokens: enforcing well-bracketed control flow and stack encapsulation using linear capabilities, *Proc. ACM Program. Lang.* 3 (POPL) (2019) 19:1–19:28. <https://doi.org/10.1145/3290332>
- [32] M. Patrignani, D. Devriese, F. Piessens, On modular and fully-abstract compilation, in: *IEEE 29th Computer Security Foundations Symposium, CSF 2016, IEEE Computer Society, 2016*, pp. 17–30. <https://doi.org/10.1109/CSF.2016.9>
- [33] V. Koutavas, M. Wand, Small bisimulations for reasoning about higher-order imperative programs, in: J.G. Morrisett, S.L.P. Jones (Eds.), *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2006, ACM, 2006*, pp. 141–152. <https://doi.org/10.1145/1111037.1111050>
- [34] D. Sangiorgi, N. Kobayashi, E. Sumii, Environmental bisimulations for higher-order languages, in: *22nd IEEE Symposium on Logic in Computer Science (LICS 2007)*, Proceedings, IEEE Computer Society, 2007, pp. 293–302. <https://doi.org/10.1109/LICS.2007.17>
- [35] V. Koutavas, P.B. Levy, E. Sumii, From applicative to environmental bisimulation, in: M.W. Mislove, J. Ouaknine (Eds.), *Twenty-seventh Conference on the Mathematical Foundations of Programming Semantics, MFPS 2011, Vol. 276 of Electronic Notes in Theoretical Computer Science*, Elsevier, 2011, pp. 215–235. <https://doi.org/10.1016/j.entcs.2011.09.023>
- [36] J. Madiot, D. Pous, D. Sangiorgi, Modular coinduction up-to for higher-order languages via first-order transition systems, *Log. Methods Comput. Sci.* 17 (3) (2021). [https://doi.org/10.46298/LMCS-17\(3:25\)2021](https://doi.org/10.46298/LMCS-17(3:25)2021)
- [37] D. Biernacki, S. Lenglet, P. Polesiuk, A complete normal-form bisimilarity for algebraic effects and handlers, in: Z.M. Ariola (Ed.), *5th International Conference on Formal Structures for Computation and Deduction, FSCD 2020, Vol. 167 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik*, 2020, pp. 7:1–7:22. <https://doi.org/10.4230/LIPIcs.FSCD.2020.7>
- [38] A.S. Murawski, N. Tzevelekos, Algorithmic games for full ground references, *Formal Methods Syst. Des.* 52 (3) (2018) 277–314. <https://doi.org/10.1007/s10703-017-0292-9>
- [39] D. Hirschhoff, I. Quémerais, D. Sangiorgi, First-order store and visibility in name-passing calculi, in: *CONCUR 2025, LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik*, 2025, pp. 23:1–23:21. <https://doi.org/10.4230/LIPIcs.CONCUR.2025.23>
- [40] I. Quémerais, G. Jaber, K. Sakayori, D. Sangiorgi, Concurrent visibility: higher-order concurrency with first-order store, in: *CONCUR 2026, LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik*, 2026.
- [41] C. Bodei, V.D. Dinh, G. Ferrari, Checking global usage of resources handled with local policies, *Sci. Comput. Program.* 133 (2017) 20–50. <https://doi.org/10.1016/J.SCICO.2016.06.005>
- [42] M. Bartoletti, P. Degano, G. Ferrari, R. Zunino, Local policies for resource usage analysis, *ACM Trans. Program. Lang. Syst.* 31 (6) (2009) 23:1–23:43. <https://doi.org/10.1145/1552309.1552313>

- [43] M. Bartoletti, P. Degano, G. Ferrari, R. Zunino, Types and effects for resource usage analysis, in: H. Seidl (Ed.), FOSSACS 2007, Lecture Notes in Computer Science, Springer, 2007, pp. 32–47. https://doi.org/10.1007/978-3-540-71389-0_4
- [44] L. Bettini, V. Bono, R. De Nicola, G. Ferrari, D. Gorla, M. Loret, E. Moggi, R. Pugliese, E. Tuosto, B. Venneri, The klaim project: theory and practice, in: Global Computing '23, Lecture Notes in Computer Science, Springer, 2003, pp. 88–150. https://doi.org/10.1007/978-3-540-40042-4_4