



HAL
open science

First-order store and visibility in name-passing calculi

Daniel Hirschhoff, Iwan Quémerais, Davide Sangiorgi

► **To cite this version:**

Daniel Hirschhoff, Iwan Quémerais, Davide Sangiorgi. First-order store and visibility in name-passing calculi. 36th International Conference on Concurrency Theory (CONCUR 2025), Aug 2025, Aarhus (Denmark), Denmark. <10.4230/LIPIcs>. <hal-05033251>

HAL Id: hal-05033251

<https://hal.science/hal-05033251v1>

Submitted on 23 Apr 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

First-order store and visibility in name-passing calculi

Daniel Hirschkoff Iwan Quémerais Davide Sangiorgi

Abstract

The π -calculus is the paradigmatical name-passing calculus. While being purely name-passing, it allows the representation of higher-order functions and store.

We study how π -calculus processes can be controlled so that computations can only involve storage of first-order values. The discipline is enforced by a type system that is based on the notion of visibility, coming from game semantics. We discuss the impact of visibility on the behavioural theory. We propose characterisations of may-testing and barbed equivalence, based on (variants of) trace equivalence and labelled bisimilarity, in the case where computation is sequential, and in the case where computation is well-bracketed.

1 Introduction

The π -calculus is one of the best known process calculi. It has a handful of operators — parallelism, input, output, restriction being the main ones — with which it achieves an amazing expressive power. Yet, it has a rich algebraic theory and a wide spectrum of proof techniques. In the π -calculus, computation is communication, by means of message-passing synchronisations between processes. Names are the communication units; new names may be created, and may themselves be exchanged in communications. When used for modelling, π -calculus is normally typed, so to be able to support structured data values, beginning with basic first-order values such as integers and booleans.

The π -calculus has been advocated as a model to interpret, and give semantics to, languages with higher-order features. For instance, the representation of functions as π -calculus processes has been widely studied, beginning with Milner’s seminal work [14]. Usually, higher-order languages include forms of *store*. It is therefore important to understand how expected semantic properties of store can be captured in a language for pure concurrency such as the π -calculus.

An aspect of store with strong semantic consequences is the distinction between higher-order and first-order store. The term ‘first-order store’ refers to a memory model where variables can hold only basic data types such as integers and booleans. This contrasts with a higher-order store, where variables can also store functions, or code, or references to these. Limiting the store to first-order data makes it easier to reason about program behaviours, and verify behavioural properties. It may also enhance security and predictability guarantees by preventing unintended behaviours associated with dynamic function storage and code

injection. A striking example of a property that may be broken with higher-order store is termination. For instance, the simply-typed λ -calculus is strongly normalising and therefore terminating. This property is maintained by the addition of first-order store, but not with higher-order store, as functions may recursively reference and invoke each other through the store. In game semantics of higher-order languages, the distinction between first-order and higher-order store has led to studying the concept of *visibility*. Intuitively, visibility ensures us that each move in a computational game is justified by prior interactions. With the introduction of higher-order store, the visibility condition has to be relaxed, reflecting the added complexity in tracking information flow within a program.

The goal of this paper is precisely to understand the meaning of first-order store and visibility within the π -calculus. In functional languages based on the λ -calculus, imposing first-order store is simply determined by the constraint that all references may only store first-order values. In the π -calculus, in contrast, there is no explicit notion of reference, and therefore the meaning of ‘first-order store’ requires some care. In π -calculi, a piece of store is usually represented by an output message. For instance $\bar{h}(5)$ may be thought of as the statement that name h currently contains 5. A process that inputs at h and then re-emits a message at h , such as $h(z). \bar{h}(z+1)$, is a process that acts on such a store. In this case, h is a first-order name, hence h implements a piece of first-order store. An example of higher-order store is given by the process

$$P \stackrel{\text{def}}{=} a(b). ((\bar{b}(d). !d. R) \mid c. \bar{b}(d'). d'. R') \quad (1)$$

Here, $\bar{b}(d)$ is a bound output, and represents the output of a ‘new’ name d , and similarly for $\bar{b}(d')$; whereas, e.g., $d. R$ and $c. Q$ represent inputs in which no value is exchanged (i.e., d and c carry **unit** values). In P , name a is higher-order, in that a carries a (pointer to a) service. In the input at a , the service name received in the input (the parameter b) is used after the input, but it is also stored, and used later when c is invoked. In other words, the visibility of c (the services that c may call when invoked) depends on interactions that have occurred earlier, at a . For this reason, P implements a form of higher-order store. When this does not happen, that is, the visibility of a higher-order name (the set of services that may be called when the name is invoked) may be determined at the point in which the name is created, the process implements a form of first-order store. (The word ‘service’ is freely used, here and elsewhere, to indicate the input end of a name; in π -calculus, a name may have multiple input occurrences, in arbitrary syntactic positions within a process, in contrast with, say, the function declarations in functional languages, which are supposed to be unique.)

The amazing expressive power of the π -calculus also makes the contexts of the calculus very discriminating; as a consequence, behavioural equivalences, which are supposed to be preserved by the contexts of the calculus, are rather demanding relations. A well-established way of imposing constraints on the set of legal contexts, thus obtaining usefully coarser behavioural equivalences, is to adopt behavioural type systems. Types may be used in order to capture communication patterns on the usage of names, such as capabilities, linearity, sessions, and so on on, e.g., [18, 10, 5, 1]. Type systems have also been designed to capture specific properties of processes, such as termination, deadlock-freedom, lock-freedom, e.g., [23, 11, 17].

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 behavioural equalities that only hold in presence of types. For instance, in the case of may testing (or contextual equivalence) this may yield to refinements or modifications of the notion of trace equivalence; whereas in the case of barbed equivalence this may lead to refinements or modifications of the standard notion of labelled bisimilarity. The resulting proof techniques must be sound with respect to initial contextually-defined relations; ideally, they should also be complete.

In this paper we propose a type system that guarantees first-order store in π -calculus processes. The type system makes use of *visibility functions*, that is, partial finite functions that specify, for each higher-order name b that may occur free in a process, the set of services (i.e., higher-order names) that may be invoked when a call at b is made. In order to better understand the behavioural effects of visibility, we combine it first with *sequentiality* and then with *well-bracketing*. ‘Sequentiality’ intuitively indicates the existence of a single thread of computation. That is, at any point there is a single process that is active and decides what the next computation step can be. In the encodings of the λ -calculus [15, 20], a process is active, i.e., it carries the thread, when it contains an unguarded output at a higher-order name. When encoding λ -calculi with first-order references, an active process may also expose the thread by exhibiting an input at a first-order name. We follow such a convention also in the paper: higher-order names carry the thread via outputs, first-order names via inputs (other choices technically would however be possible). Sequentiality allows us to expose, in isolation, the causality properties of service calls, within a standard (interleaving) semantics. These aspects would be partially or totally obscured by the presence of other concurrent threads.

In higher-order languages, the effects of having first-order store are particularly strong in absence of control operators; that is, using a terminology borrowed from game semantics, in a ‘well-bracketed’ setting, where the call-return interaction behaviour between a term and its context follows a stack discipline. We therefore also consider the refinement of our visibility-based type system under the well-bracketing hypothesis. To see an example of the effect of visibility, consider the following well-bracketed example:

$$\nu h (\bar{h}\langle 0 \rangle \mid \bar{a}(c, q). (!c. h(z)). \bar{h}\langle z + 1 \rangle \mid q(y). Q)$$

Here, an external service a is called, exporting a local service c capable of incrementing a private reference h , and with a return channel q . If visibility holds, then the external environment cannot store c . Therefore, after providing an answer at q , the environment is unable to call c again, thus incrementing h , unless an access to c is explicitly provided again (by process Q). This property may be used to perform optimisations, possibly including garbage-collection of the service c if at some point c is not used anymore in Q .

We then study characterisations for both may testing and barbed equivalence, as special forms of trace equivalence and of labelled bisimilarity. Following the tradition of testing, the language includes special success signals, that are used by testers (i.e, processes running in parallel with the testees), and may be used at any time, to report success.

When developing such characterisations, some care is necessary so to distinguish between the visibility of the tested and of the tester processes, as they need not be the same. For this,

we introduce visibility-constrained LTSs, in which process transitions are constrained by a visibility function, intuitively expressing the possibility for processes to perform a certain transition given a certain visibility function for the external observer. As a consequence of the transition, the visibility function may need to be updated, so to yield the visibility function for the next transition. In the case of well-bracketing, the constraints given by visibility have to be coupled with those given by the stack names used to track well-bracketing (and again, a distinction has to be made between the stacks for the testers and those for the testees). Based on these LTSs, we introduce forms of trace equivalence and of labelled bisimilarity that are proved to be sound and complete, for may testing and barbed equivalence, both in the case of sequentiality and in that of well-bracketing.

We present the version of the π -calculus with higher-order and first-order names in Section 2. In Section 3, we define the type system implementing visibility for sequential processes. We define typed labelled transitions and present labelled characterisations of behavioural equivalences in Section 4. Section 5 is devoted to the extension with well-bracketing.

2 Background: the π -calculus language

The π -calculus language that will be used is called, technically, Internal π (π I) [21] (all channels exchanged are fresh – no free object outputs), with the addition of *first-order values* (such as integers and booleans), and expressions that evaluate to first-order values, called *first-order expressions*. We do not specify what exactly first-order values and first-order expressions are. We simply assume that, in a closed process, a first-order expression evaluates to a first-order value. Moreover, only the output capability of names may be exported. This means that, for instance, in an input $a(b).P$ name b may only be used in output in P . This constraint is often followed in theory and applications of the π -calculus, as it simplifies various technical matters.

We assume that names are distinguished (partitioned) into *higher-order* names and *first-order* names, and *success names*. Higher-order names are the ordinary π I names; first-order names may only carry first-order values, and they may not be passed around (the second constraint simplifies the technical details, though it is not mandatory, see Section 3). In later sections, first-order names will be used to represent references. The fact that these names may only carry first-order values, together with constraints on visibility, will ensure that the calculus does not have (implicit or explicit forms of) higher-order store. The *success* names are used, in the tradition of testing equivalences, by tester processes so to report results of experiments on tested processes (thus, the processes compared in the examples of the paper do not contain success names).

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* [13], in which names are partitioned into a collection of *types* (or *sorts*), and a sorting function maps types onto types. We assume that there is a sorting system under which all processes we manipulate are well-typed. Thus, the sorting system separates higher-order and first-order

$[Higher\text{-}order\ names, HO]$	a, b, c	$[First\text{-}order\ expressions]$	e, f
$[First\text{-}order\ names, FO]$	h, k	$[First\text{-}order\ variables]$	x, y
$[Success\ names]$	ω	$[First\text{-}order\ values]$	v, w
$[Inputs]$	$\alpha ::= a(x, b) \mid h(x)$	$[Outputs]$	$\bar{\alpha} ::= \bar{a}\langle v \rangle(b) \mid \bar{h}\langle e \rangle \mid \omega$
$[Processes]$	$P, Q, R ::= \alpha.P \mid \bar{\alpha}.P \mid !\alpha.P \mid P \mid Q \mid \nu a P \mid \nu h P$ $\mid \mathbf{0} \mid P + Q \mid \text{if } e = f \text{ then } P \text{ else } Q$		

Table 1: The syntax of the calculus

names; moreover we assume that first-order names are monadic, and that higher-order names carry a pair of a first-order value and a higher-order name. Thus $\bar{a}\langle v \rangle(b).P$ is the output of the first-order value v and of the fresh higher-order name b at the (higher-order) name a , with continuation P . When writing examples, we allow different sortings; for instance, writing $a.P$ and $\bar{b}.Q$ for inputs and outputs in which no value is exchanged.

The syntactic categories of the calculus are presented in Table 1. As usual, input and restriction are binding constructs, and give rise to the expected definitions of free and bound names. An expression or a process is *closed* if it does not have free (first-order) variables. It is intended that transitions are defined on closed processes.

The definition of structural congruence, written \equiv , and of the strong and weak (untyped) labelled transitions, written $\xrightarrow{\mu}$, \Longrightarrow , and $\xRightarrow{\hat{\mu}}$, are standard and are given in Appendix A (later we will introduce typed variants of the LTS). We note $\text{fn}(P)$ (resp. $\text{fn}(\mu)$) the set of free names of P (resp. μ). We sometimes abbreviate reductions $P \xrightarrow{\tau} P'$ as $P \longrightarrow P'$.

As behavioural equivalences we consider both may-testing and barbed equivalence. Both relations are contextual, in that they are defined by requiring a closure under all ‘observers’. They both use ‘barbs’ (or ‘success’) signals; barbed equivalence, in addition, makes use of a bisimulation game on reductions. We write \Downarrow_ω to indicate the possibility of exhibiting an ω barb, possibly after some reductions; i.e., $P \Downarrow_\omega$ holds if there is P' with $P \Longrightarrow P'$ and P' has an unguarded occurrence of ω . Thus intuitively two processes are behaviourally equal if they can produce ‘the same barbs’ under all legal observers running in parallel with them. As we are in a typed setting, observer and tested processes should have compatible typings. The meaning of ‘compatible’ depends on the specific type system. For visibility, the meaning will be specified in Sections 3 and 5. We write $\Delta \vdash P$, and say that P is a Δ -process, if P is typable under Δ , and $\Delta \vdash P, Q$ if both $\Delta \vdash P$ and $\Delta \vdash Q$ hold. Behavioural equivalences are defined on closed processes; they are extended to open processes (i.e., processes with free first-order variables) in the usual manner, by considering all possible closing substitutions. We only consider *weak* behavioural equivalences (i.e., equivalences that abstract from internal moves of processes, the reductions), as these are, in practice, the most useful ones.

Definition 1 (May testing). Suppose that Γ, Δ are compatible. Then two Δ -processes P and Q are *may testing equivalent at Γ* , written $P \cong_\Delta^\Gamma Q$, if for all R -processes and success name ω , we have $(R \mid P) \Downarrow_\omega$ iff $(R \mid Q) \Downarrow_\omega$.

Definition 2 (Barbed bisimilarity and equivalence). *Barbed bisimilarity* is the largest symmetric relation \approx s.t. $P \approx Q$ implies:

1. whenever $P \longrightarrow P'$, there exists Q' such that $Q \Longrightarrow Q'$ and $P' \approx Q'$;
2. for each ω , if $P \Downarrow_\omega$ then also $Q \Downarrow_\omega$.

Suppose that Γ, Δ are compatible. Then two Δ -processes P and Q are *barbed equivalent at Γ* , written $P \approx_\Delta^\Gamma Q$, if for all R with $\Gamma \vdash R$ it holds that $P \mid R \approx Q \mid R$.

We write $P \cong Q$ (resp. $P \cong Q$) when $P \approx_\Delta^\Gamma Q$ (resp. $P \approx_\Delta^\Gamma Q$) holds under any typing Δ for P, Q and any Γ compatible with Δ .

We now introduce some notations that will be useful below. In the paper we use partial functions whose codomains are finite powersets. If Δ is such a function, then $\Delta; i \mapsto A$ is the extension of Δ in which i is mapped onto A , assuming that i is not in $\text{dom}(\Delta)$; conversely Δ^{-i} is the restriction of Δ in which the entry for i is erased, both in the domain and codomain of Δ . We also write $i \mapsto A, j$ as an abbreviation for $i \mapsto (A \cup \{j\})$; and $\Delta \downarrow i$ (resp. $\Delta \not\downarrow i$) means that Δ is (resp. is not) defined on i . Function $\Delta_1 \cap \Delta_2$ has domain $\text{dom}(\Delta_1) \cap \text{dom}(\Delta_2)$, and then is defined as $(\Delta_1 \cap \Delta_2)(p) = \Delta_1(p) \cap \Delta_2(p)$; and likewise for $\Delta_1 \cup \Delta_2$. We write $\Delta_1 \subseteq \Delta_2$ if $\text{dom}(\Delta_1) \subseteq \text{dom}(\Delta_2)$, and for each $i \in \text{dom}(\Delta_1)$, we have $\Delta_1(i) \subseteq \Delta_2(i)$.

3 A Type System for Visibility

Our type system for visibility is built on top of that for sequentiality [3]. While visibility could be defined as a standalone notion, the combination with sequentiality makes the behavioural effects sharper and easier to grasp.

Intuitively, sequentiality ensures us that at any time at most one interaction can occur; i.e., there is a single computation thread. A process that holds the thread, called *active*, decides what the next interaction can be. It does so by offering a single particle (input or output) that controls the thread. A given name may exercise the control on the thread either in output or in input. Following [3], we require that, in higher-order names, the thread is carried by outputs, whereas, in first-order names, it is carried by inputs. An input that carries the thread maintains the thread, whereas an output releases it. (These appear to be, practically, the most interesting combinations; other combinations could be adopted, with the expected modifications.) Thus, for instance, the following process P is well-typed:

$$P \stackrel{\text{def}}{=} (h(x). \bar{a}\langle x \rangle(b). (b(z, c'). Q \mid \bar{h}\langle x \rangle)) \mid \bar{k}\langle v \rangle$$

The initial particles in P are the input at h and the output at k ; however only the input at h carries the thread, as h and k are first-order. Underneath the input, the thread is maintained and released in the output at a (the thread will be acquired by the process in the external environment performing the input at a); later, when the input at b is performed, process Q obtains the thread back. Sequentiality however does not exclude non-determinism: e.g., an output particle may have the possibility of interacting with different inputs.

The type system for *visibility* intuitively allows us to specify which possible higher-order outputs can be performed by a process that becomes active. Type environments are partial functions, called *visibility functions*, acting on higher-order names and, for active processes, on the special token $*$, representing the active thread. (Later, when studying well-bracketing, type environments will have both a visibility function and a name stack.)

Definition 3 (Visibility function). A *visibility function* is a finite partial function $\Delta : \mathsf{HO} \cup \{*\} \rightarrow \mathcal{P}_{\text{fin}}(\mathsf{HO})$. Moreover the function should be *transitive*, that is, for all $o \in \text{dom}(\Delta)$, if $\Delta \downarrow a$ and $a \in \Delta(o)$, then $\Delta(a) \subseteq \Delta(o)$.

To see the use of $*$, suppose $\Delta(*) = \{a, b\}$; this means that a Δ -process is active, and may only perform outputs at a and b , as in the following process:

$$\bar{a}\langle v \rangle(c).c(d).\bar{b}\langle w \rangle.\mathbf{0}$$

Name c is created in the output at a and therefore c inherits a 's visibility.

Similarly, if $\Delta(d) = \{a, b\}$ then, in a Δ -process $d(x, c).P$, the continuation P of the input at d may only perform outputs at a , b or at c (as c is the higher-order name received in the input). The use of visibility functions also allows us to sort out implicit form of higher order store; in the following example, p, q are higher order names used for continuations.

Example 4 (Implicit forms of higher-order store). In the λ -calculus, if a term without higher-order references is well-bracketed then it necessarily obeys visibility. In our calculus, even though we do not have explicit higher-order references, we can simulate higher-order store. Consider

$$P \stackrel{\text{def}}{=} \bar{a}\langle b, p \rangle.b(n, c, q).\bar{q}\langle n \rangle.p(m).\bar{c}.\mathbf{0}$$

In P the visibility $\Delta(p)$ is defined when the output at a is made, therefore $c \notin \Delta(p)$, and the final output at c breaks visibility; indeed, the input at b , which justifies the output at c , is answered by the output at q before the output at c .

Transitivity ensures us that the visibility functions, viewed as ‘named sets’, represent transitive sets. Transitivity is needed to guarantee the behavioural expectations of visibility. If a process P calls a service a , with actual (higher-order) parameter b , then an output triggered by such a call (and possibly directed back to P) should be either at b or at a name in the visibility of a ; however, the server a might call an external service c ; this call should not allow an increase of the visibility associated to a ; hence the transitivity constraint that the visibility of c is contained in that of a . Technically, transitivity is necessary to ensure preservation of typing under reduction (see Theorem 10 and Example 11).

Remark 5 (Transitivity and functional languages). The issue of transitivity does not arise in the semantics of purely functional languages such as the λ -calculus [6]: visibility functions do appear, but the domain and codomain of such functions are disjoint (technically, in the game semantics terminology, the domain contains names introduced by Player, whereas the codomain contains names introduced by Opponent). This makes transitivity trivial.

$$\begin{array}{c}
\text{O-HO} \frac{a \in \Delta(*) \quad \Delta^{-*}; b \mapsto \Delta(*), b \vdash P}{\Delta \vdash \bar{a}\langle e \rangle(b). P} \qquad \text{I-HO} \frac{\Delta; * \mapsto \Delta(a), b \vdash P}{\Delta \vdash a(x, b). P} \\
\\
\text{O-FO} \frac{\Delta \vdash P \quad \Delta \not\vdash *}{\Delta \vdash \bar{h}\langle e \rangle. P} \qquad \text{I-FO} \frac{\Delta \vdash P \quad \Delta \downarrow *}{\Delta \vdash h(x). P} \qquad \text{REP} \frac{\Delta \vdash \alpha. P \quad \Delta \not\vdash *}{\Delta \vdash !\alpha. P} \qquad \text{SUCC} \frac{}{\Delta \vdash \omega. P} \\
\\
\text{RES-HO} \frac{\Delta \vdash P}{\Delta^{-a} \vdash \nu a P} \qquad \text{RES-FO} \frac{\Delta \vdash P}{\Delta \vdash \nu h P} \qquad \text{NIL} \frac{\Delta \not\vdash *}{\Delta \vdash \mathbf{0}} \qquad \text{SUM} \frac{\Delta \vdash P_1 \quad \Delta \vdash P_2}{\Delta \vdash P_1 + P_2} \\
\\
\text{PAR} \frac{\Delta_1 \vdash P_1 \quad \Delta_2 \vdash P_2 \quad (\Delta_1, \Delta_2) \in \text{split}(\Delta)}{\Delta \vdash P_1 \mid P_2} \qquad \text{IF} \frac{\Delta \vdash P_1 \quad \Delta \vdash P_2}{\Delta \vdash \text{if } e = f \text{ then } P_1 \text{ else } P_2}
\end{array}$$

Figure 1: Typing rules for visibility

In the typing of a parallel composition, the visibility function has to be split between the two components of the composition. This is important for elements of the domain of the function that are linear, and therefore may only appear in one, but not both, components; this is the case for $*$, and the linear continuation names that will be introduced in Section 5.

Definition 6 (Split). Let Δ be a visibility function, we define $\text{split}(\Delta)$ as the set of pairs (Δ_1, Δ_2) such that $\Delta_1 \cup \Delta_2 = \Delta$ and $* \notin \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2)$.

The typing rules are given in Table 1. We comment on a few of them. Rule **O-HO** specifies that an output at a **HO** name a owns the thread, and that name a must be in the current visibility. In the continuation P (which is not active) the visibility for b is determined by the current visibility $\Delta(*)$ together with b itself, to allow for recursive calls at b . Symmetrically, in rule **I-HO**, a process performing an input at a acquires the thread, with a visibility given by $\Delta(a)$ plus the received name b . A replicated input is well-typed only if it is not active (rule **REP**). When typing parallel composition, we use **split** to insure that at most one of the components holds the thread. We simply remove the name from Δ when typing a restriction (rule **RES-HO**). Visibility plays no role in rules **I-FO** and **O-FO**; the only checks made are given by sequentiality, concerning the respect of the thread.

Example 7. Consider

$$\begin{aligned}
P \stackrel{\text{def}}{=} & a(x, b). \bar{b}\langle x + 1 \rangle(d). a(x', b'). \bar{b}\langle x + x' \rangle(d'). \mathbf{0} \\
& \mid a(x_1, b_1). \bar{b}_1\langle x_1 \rangle(d_1). d_1(d_2). \bar{c}\langle x + 1 \rangle(c_1). \mathbf{0}
\end{aligned}$$

The process has three inputs at a ; the first two are sequential, and make no requirements on visibility, as their continuations only perform outputs at names received in the input; the third one, in addition, has the effect of creating a name d_1 , whose input may activate an output at the free name c . This input requires that in the visibility Δ of P , we find $c \in \Delta(a)$.

Example 8. (In this example, we ignore the first-order components of actions, as they are not relevant.) A process of the form

$$P \stackrel{\text{def}}{=} \bar{a}(b).(Q \mid c(d).\bar{b}(d').R)$$

is not typable: the initial visibility of c may not include b , which is created fresh by a ; yet, P may perform an output at b after the input at c . Indeed, while P does not directly create a form of higher-order store, it may do so indirectly: the external environment, after completing the treatment of the initial call at a , is supposed to ‘forget’ the higher-order name b . However, this is not the case, as the environment may regain access to b by invoking the (unrelated) name c . In contrast P would become typable (in the straightforward extension of our type system with polyadicity) if the initial output at a would be replaced by $\bar{a}(b, c)$: now c becomes a fresh name, initially unknown, and created in the initial output, together with b , so that the visibility for b and c can be the same and can include both names.

The following lemma shows that a visibility function may always be enlarged (both in the domain and codomain), as long as sequentiality (i.e., being defined on $*$) is respected.

Lemma 9 (Weakening). *Suppose $\Delta_1 \vdash P$, $\Delta_1 \subseteq \Delta_2$, and $(\Delta_2 \downarrow * \text{ iff } \Delta_1 \downarrow *)$. Then $\Delta_2 \vdash P$.*

Some other technical lemmas that are used in proofs are presented in Appendix B.1.

Theorem 10 (Subject reduction). *Suppose $\Delta \vdash P$ and $P \xrightarrow{\mu} P'$. We have:*

- *if μ is a τ action or a first-order input, then $\Delta \vdash P'$; the same holds if μ is a first-order output and $\Delta \not\downarrow *$;*
- *if $\mu = \bar{a}(v)(b)$ then $a \in \Delta(*)$ and $\Delta^{-*}; b \mapsto \Delta(*), b \vdash P'$;*
- *if $\mu = a(v)(b)$ and $\Delta \not\downarrow *$, then $\Delta; * \mapsto \Delta(a), b \vdash P'$.*

Example 11. Transitivity of visibility functions is necessary. Indeed, consider process $P \stackrel{\text{def}}{=} a.\bar{b}.Q \mid \bar{a}$, and the *non-transitive* function $\Delta = a \mapsto \{a, b\}, * \mapsto a$. We have $\Delta \vdash P$. However $P \xrightarrow{\tau} \bar{b}.Q$, with $\Delta \not\downarrow \bar{b}.Q$, violating Subject Reduction (Theorem 10).

To adapt the generic definitions of behavioural equivalence in Section 2 to the type system for visibility, we only have to specify the meaning of compatibility between typing environments (which in our case are visibility functions). Compatibility means ensuring that processes typed under the two typing environments cannot both be active.

Definition 12. Two typings Δ_1 and Δ_2 are *compatible* if $* \notin \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2)$.

We discuss some examples of equivalences in which the behavioural constraints imposed on legal observers by visibility are essential: the equivalences would otherwise be broken.

Example 13. Consider a process

$$P \stackrel{\text{def}}{=} \bar{a}(b, c). P' \quad \text{for} \quad P' \stackrel{\text{def}}{=} (b(d). R) \mid c(c_1). \bar{a}(b', c'). c'(c_2). \bar{b}(d'). Q$$

where we assume b does not occur in Q and R . Process P interrogates an external service a , providing access to local services b and c , and the derivative is P' . Now, P' may either be interrogated at b , or at c . In the former case, the output $\bar{b}(d'). Q$ becomes dead code (as the environment may not install any input at b). In the latter case, the environment is again called at a ; now, since visibility ensures us the environment cannot remember b , the input and the output at b in P' may only interact with each other; we can therefore inline the call at b . We can express this behaviour using expansion as follows:

$$P \cong \bar{a}(b, c). \left(\begin{array}{l} b(d). (R \mid c(c_1). \bar{a}(b', c'). c'(c_2). \mathbf{0}_*) \\ + \quad c(c_1). \bar{a}(b', c'). c'(c_2). \tau. \nu d' (Q \mid R\{d'/d\}) \end{array} \right)$$

where $\mathbf{0}_*$ indicates an active process that cannot perform any action (e.g., $\nu b_2(\bar{b}_2. \mathbf{0})$). (This equality holds under any visibility, hence we use the symbol \cong .)

If we replaced the second output at a with an output at a different free name, say a' , then the processes could still be typable, and the equalities above would still hold. However, the equalities would not hold if we replaced the second output at a with an output at c_1 . Indeed, while the environment may not store b or c for using them in a later call at a (or at an unrelated name such as a'), the environment is allowed to use b and c within services that are created following the first call at a ; for instance, the environment could be

$$a(b, c). \bar{c}(c_1). !c_1(c'_1). \bar{b}(d''). \mathbf{0}$$

The following example illustrates the use of visibility to perform garbage-collection on names that will never be used. We refer to Appendix B.3 for the full definition of the corresponding pruning function.

Example 14. Consider a typing environment Δ , a Δ -process

$$P \stackrel{\text{def}}{=} !a. Q_1 \mid !b. \bar{c}. !a. Q_2 \mid !d. \bar{a}. \mathbf{0}$$

and a visibility Γ compatible with Δ such that $\Gamma \cup \Delta$ is transitive. Γ is intended here to represent the observer's visibility. Suppose $a \notin \Gamma(*)$; that is, the observer is active but its current visibility does not include a (which may however be in the visibility of other names in Γ). We then have $P \cong_{\Gamma}^{\Delta} !b. \bar{c}. \mathbf{0}$. The pruning is possible because, intuitively, the transitivity of $\Delta \cup \Gamma$ ensures us that an observer, whenever active, cannot have a in its current visibility. For instance, if the observer interrogates b , then the transitivity of the visibility functions ensures us that $a \notin \Delta(b)$; then $c \in \Delta(b)$ (which is necessary for P to be typable) gives us also $a \notin \Gamma(c)$. Similarly, exploiting transitivity, we deduce that also $d \notin \Gamma(*)$, and that an observer may never be able to interrogate d ; thus we can also remove the input on d .

$$\begin{array}{ccc}
\text{TrO-HO} \frac{P \xrightarrow{\bar{a}\langle v \rangle(b)} P' \quad a \in \text{dom}(\Gamma) \quad \Gamma \not\downarrow *}{[\Gamma] P \xrightarrow{\bar{a}\langle v \rangle(b)} [\Gamma; * \mapsto \Gamma(a), b] P'} & & \text{TrI-HO} \frac{P \xrightarrow{a\langle v \rangle(b)} P' \quad a \in \Gamma(*)}{[\Gamma] P \xrightarrow{a\langle v \rangle(b)} [\Gamma^{-*}; b \mapsto \Gamma(*), b] P'} \\
\\
\text{TrO-FO} \frac{P \xrightarrow{\bar{h}\langle v \rangle} P' \quad \Gamma \downarrow *}{[\Gamma] P \xrightarrow{\bar{h}\langle v \rangle} [\Gamma] P'} & & \text{TrI-FO} \frac{P \xrightarrow{h\langle v \rangle} P' \quad \Gamma \not\downarrow *}{[\Gamma] P \xrightarrow{h\langle v \rangle} [\Gamma] P'} & & \text{TrTau} \frac{P \xrightarrow{\tau} P' \quad \Gamma \not\downarrow *}{[\Gamma] P \xrightarrow{\tau} [\Gamma] P'}
\end{array}$$

Figure 2: Visible LTS

4 Labelled Semantics

4.1 Trace-based proof techniques

The definition of may-testing makes use of contextual quantifications. We study here proof techniques without such quantifications, as forms of trace equivalence. For lack of space, analogous proof techniques for barbed equivalence are only reported in Appendix C. The first step is to express the transitions for a process that are allowed by an observer subject to certain visibility constraints. We call these the *type-allowed transitions*. They are of the form $[\Gamma] P \xrightarrow{\mu} [\Gamma'] P'$, to be read ‘an observer that has visibility Γ allows process P to make a transition μ and, as a result of the transition, the process becomes P' and the visibility of the observer becomes Γ' ’. Type-allowed transitions are defined by the rules in Figure 2. Intuitively, an observer typable under Γ can initially test inputs of processes at names in $\Gamma(*)$, if $\Gamma \downarrow *$, and outputs of processes in $\text{dom}(\Gamma)$ otherwise. The weak transitions \Longrightarrow and $\xRightarrow{\mu}$, are defined in the usual manner from the strong one; for instance, $[\Gamma] P \xRightarrow{\mu} [\Gamma'] P'$ holds if $[\Gamma] P \Longrightarrow [\Gamma] P''$, and $[\Gamma] P'' \xrightarrow{\mu} [\Gamma'] P'''$, and $[\Gamma'] P''' \Longrightarrow [\Gamma'] P'$, for some P'' and P''' .

An action τ is called a *reduction* (or an *internal action*); the remaining actions are called *interactions*. We use \mathcal{T} to range over *traces*, which are finite sequences of interactions $\mu_1 \dots \mu_n$ ($n \geq 0$) in which the bound names are all distinct and different from the free names (the definition of bound and free names is as for processes, viewing a trace as a sequence of prefixes). A trace $\mu_1 \dots \mu_n$ is *trace of P_0 under visibility Γ_0* if there are P_i, Γ_i with $[\Gamma_i] P_i \xrightarrow{\mu_i} [\Gamma_{i+1}] P_{i+1}$, for $0 \leq i < n$.

Definition 15 (Typed trace equivalence). Two Δ -processes P and Q are *trace equivalent at Γ* , written $P \simeq_{\Delta}^{\Gamma} Q$, if Δ, Γ are compatible and P, Q have the same sets of traces under Γ .

We show the main compositionality lemmas, and a trace-based characterisation of may testing. Other technical lemmas needed in proofs may be found in Appendix B.2.

Lemma 16 (Restriction). *If $P \simeq_{\Delta}^{\Gamma} Q$ then also $\nu a P \simeq_{\Delta - a}^{\Gamma - a} \nu a Q$.*

Lemma 17 (Input). *If $P \simeq_{\Delta}^{\Gamma} Q$ and $\Gamma \not\downarrow *$; then also $a(x, b). P \simeq_{\Delta - *, b}^{(\Gamma; * \mapsto \Gamma(b))^{-b}} a(x, b). Q$.*

Lemma 18 (Output). *If $P \simeq_{\Delta}^{\Gamma} Q$ and $\Gamma \not\downarrow *$, then also $\bar{a}\langle e \rangle(b). P \simeq_{(\Delta; * \mapsto \Delta(b))^{-b}}^{\Gamma^{-*}, b} \bar{a}\langle e \rangle(b). Q$.*

Lemma 19 (Parallel composition). *Suppose $P \simeq_{\Delta}^{\Gamma} Q$ and $(\Gamma_1, \Gamma_2) \in \mathbf{split}(\Gamma)$. Then $\Gamma_2 \vdash R$ implies $P \mid R \simeq_{(\Delta \cup \Gamma_2)^+}^{\Gamma_1} Q \mid R$ where $(-)^+$ indicates the transitive closure.*

In the proof of Lemma 19, a trace of $P \mid R$ is split into a trace for P and one for R . We have then to ensure that the condition of the lemma are maintained throughout the trace, notwithstanding modifications of the visibility functions. Some care is needed in case of synchronisations between P and R , where the typing of the external observer has to remain the same; we also use the restriction Lemma 16. Other cases also exploit the extension Lemma 36, and the transitivity of the visibility function, to make sure that the views on free names are not increased along the trace.

Theorem 20 (Characterisation of may testing using trace equivalence). *Suppose $\Delta \vdash P, Q$ then $P \simeq_{\Delta}^{\Gamma} Q$ if and only if $P \cong_{\Delta}^{\Gamma} Q$.*

The sketch of the proofs is the same as for characterisations of may testing with trace equivalence in untyped calculi, exploiting substitutivity properties of trace equivalence in the case of Soundness, and reason by contradiction in the case of Completeness. In addition, one has to take care of the observer visibilities, how it changes, and use type-allowed transitions. A similar results of characterisation also hold for barbed equivalence, using a typed form labelled bisimilarity, and are presented in Appendix C.

We can use typed trace equivalence to prove equalities such as those in Examples 13 and 14; the proofs are simple, following the definition of the equivalence and exploiting type-allowed transitions. See also Appendix B.3 for the proof about the pruning function that generalises Example 14. (Similarly, the proofs can be carried out for barbed equivalence, using the labelled bisimilarity in Appendix C.)

4.2 Visibility property on traces

We formalise the link between our type system and the notion of visibility in game semantics [6]. The reader not familiar with game semantics may safely skip the section. In game semantics, the notions of visibility and well-bracketing are defined based on a justification relation between moves performed by Player. These moves correspond to interactions on higher-order names in our setting, and the notion of justification can be formulated as follows.

Definition 21. Let μ_1, μ_2 be two interactions at HO names. We say that μ_1 *justifies* μ_2 , noted $\mu_1 \curvearrowright \mu_2$, if: (i) either $\mu_1 = a\langle v \rangle(b)$ and $\mu_2 = \bar{b}\langle v' \rangle(c)$, or (ii) $\mu_1 = \bar{a}\langle v \rangle(b)$ and $\mu_2 = b\langle v' \rangle(c)$. When this is not the case, we write $\mu_1 \not\curvearrowright \mu_2$.

$\mu_1 \curvearrowright \mu_2$ captures the dependency induced by the fact that μ_2 performs an interaction at a name that has been previously introduced by μ_1 . In contrast with the situation in game semantics, though, there is no notion of question and answer in our setting.

We can now define the notion of view. In the definition below, we allow a starting visibility A , because we want to consider processes that already have some knowledge about the observer (Opponent).

Definition 22 (View). Given $A \subseteq \text{HO}$, we define the *view with starting visibility A of a trace \mathcal{T}* , written $\text{view}_A(\mathcal{T})$, as follows:

$$\begin{aligned} \text{view}_A(\varepsilon) &= A & \text{view}_A(\mathcal{T}_1\mu_1\mathcal{T}_2\mu_2) &= \text{view}_A(\mathcal{T}_1) \cup \text{nHO}(\mu_1) \cup \text{nHO}(\mu_2) \text{ if } \mu_1 \curvearrowright \mu_2 \\ \text{view}_A(\mu_0 \dots \mu_n) &= A \cup \text{nHO}(\mu_n) \text{ if } \forall 0 \leq i \leq n-1, \mu_i \not\curvearrowright \mu_n \end{aligned}$$

where $\text{nHO}(\mu)$ denotes the set of higher-order names appearing in the action μ .

Views capture the idea that when an action is justified by another one, the names created in between must have been forgotten, otherwise this would mean that the process had the ability to store these names in order to reuse them. In the last clause in Definition 22, the action is justified by an action that happened before the beginning of the trace, therefore it has the starting visibility. In particular, $\text{view}_\emptyset(\mathcal{T})$ yields the usual definition in game semantics.

Definition 23 (Respecting views). Consider $A \subseteq \text{HO}$, a trace $(\mu_0 \dots \mu_n)$ *respect views under starting visibility A* if $\text{fn}(\mu_0) \subseteq A$ and $\forall 1 \leq i \leq n, \text{fn}(\mu_i) \subseteq \text{view}_A(\mu_0 \dots \mu_{i-1})$

Lemma 24. *Consider Δ, Γ, P such that $\Delta \vdash P$, Γ is compatible with Δ , and $\Delta \cup \Gamma$ is transitive. Let $A = (\Gamma \cup \Delta)(*)$, then every trace of P under visibility Γ respects views with starting visibility A .*

The idea of the proof is that for a trace $(\mu_0 \dots \mu_n)$ there exist $P_0, \dots, P_{n+1}, \Delta_0, \dots, \Delta_n, \Gamma_0, \dots, \Gamma_{n+1}$ such that $P_0 = P, \Delta_0 = \Delta, \Gamma_0 = \Gamma$ and for all $0 \leq i \leq n, [\Gamma_i] P_i \xrightarrow{\mu_i} [\Gamma_{i+1}] P_{i+1}$ and $\Delta_i \vdash P_i$. We show by induction that for each $1 \leq i \leq n, (\Gamma_i \cup \Delta_i)(*) \subseteq \text{view}_A(\mu_0 \dots \mu_{i-1})$; it thus follows that the trace respects views under starting visibility A . The condition of $\Delta \cup \Gamma$ being transitive is crucial, otherwise $(\Gamma_i \cup \Delta_i)(*)$ could grow in unwanted ways.

5 Adding Well-Bracketing

5.1 Type system

In the previous sections we have studied visibility in connection with sequentiality. In this section we do the same for a refinement of sequentiality, namely well-bracketing. We begin by reviewing well-bracketing, following [3], and then outline how visibility is accommodated.

From sequentiality to well-bracketing. In languages without control operators, well-bracketing intuitively means that return-call interactions among terms follow a stack-based discipline. In a well-bracketed system, a service that is interrogated, say S , acquires the thread and is supposed to return a final result (unless the computation diverges) thus releasing the thread. During its computation, S may however interrogate another service, say T , which, upon completion of its computation, will return the result to S ; service S must receive the answer from T before returning the final result. The implementation of this policy requires *continuation names*, that are used to handle the return messages. When a service is

interrogated, it receives a fresh continuation name that is used once, in output, to deliver the final result. Accordingly, a client that has interrogated the service will use the continuation name in input to receive the result. Thus continuation names are *linear* [10] — they may only be used once. In contrast with [3], we do not require *receptiveness* [22] — the input-end of the name needs not be made available as soon as the name is created. Enforcing only linearity allows us to gain flexibility. To ensure the well-bracketing policy, the type system in [19] makes use of stacks.

Definition 25 ([19]). A stack, σ , is a sequence of input- and output- tagged continuation names, in which the input and output tags alternate:

$$\sigma ::= \sigma_0 \mid \sigma_I \qquad \sigma_0 ::= p^0, \sigma_I \mid \emptyset \qquad \sigma_I ::= p^I, \sigma_0 \mid \emptyset$$

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.

We write $p^0 \in \sigma$ (resp. $p^I \in \sigma$) if there is an occurrence of p^0 (resp. p^I) in σ . We write $p \in \sigma$ if either $p^0 \in \sigma$ or $p^I \in \sigma$.

Intuitively, a stack expresses the expected usage of the free continuation names in a process. For instance, if P can be typed using stack $p_1 : \circ, p_2 : \mathbf{i}, p_3 : \circ, p_4 : \mathbf{i}$, then p_1, \dots, p_4 are the free continuation names in P ; among these, p_1 will be used first, in an 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 , so it shall perform a reduction at p_3 ; the computation for P terminates with an output at p_4 . This behaviour 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.

Visibility. We now discuss how to combine the visibility in Section 3 with well-bracketing. Higher-order names are partitioned into the set **SER** of *server names*, and the set **CON** of *continuation names*. Now a, b, c range over server names, while p, q, r range over continuation names. By convention, server names carry a triple consisting of a first-order value, a server name and a continuation; whereas continuations carry a first-order value and a server name. Accordingly, outputs on higher-order names are of the form $\bar{a}\langle n \rangle(b, p)$ and $\bar{p}\langle n \rangle(b)$, and similarly for inputs. First-order names are as before.

Visibility functions are then transitive partial functions from $\mathbf{SER} \cup \mathbf{CON} \cup \{*\}$ to $\mathcal{P}_{\text{fin}}(\mathbf{SER} \cup \mathbf{CON})$. Server names are handled essentially like **HO** names in Section 3. In contrast, some care is needed on continuation names because they are linear. For this reason, some of the notations for handling visibility functions have to be revisited. A visibility function Δ is *duplicable*, written $\Delta \ddagger$, if $\text{dom}(\Delta) \subseteq \mathbf{SER}$. This predicate is useful to insure that continuation names are used linearly. In particular, the definition of splitting is refined as follows: $\text{split}(\Delta)$ is the set of pairs (Δ_1, Δ_2) such that $\Delta_1 \cup \Delta_2 = \Delta$ and $(\Delta_1 \cap \Delta_2) \ddagger$.

Typing environments include a stack; thus a typing environment is a pair $\langle \Delta \mid \sigma \rangle$ of a visibility function and a stack. In the typing rule for parallel composition, we rely on splitting as well as on $\text{inter}(\sigma_1, \sigma_2)$, the interleaving of stacks σ_1 and σ_2 .

$$\begin{array}{c}
\text{0-ser} \frac{a \in A \quad \langle \Delta; b \mapsto A, b; p \mapsto A, b \mid p^{\mathbb{I}}, \sigma \rangle \vdash P}{\langle \Delta; * \mapsto A \mid \sigma \rangle \vdash \bar{a}(e)(b, p). P} \quad \text{0-con} \frac{p \in A \quad \langle \Delta; b \mapsto A, b \mid \sigma \rangle \vdash P}{\langle \Delta; * \mapsto A \mid p^{\mathbb{O}}, \sigma \rangle \vdash \bar{p}(e)(b). P} \\
\text{I-ser}_1 \frac{\langle \Delta; a \mapsto A; * \mapsto A, b, p \mid p^{\mathbb{O}}, \sigma \rangle \vdash P}{\langle \Delta; a \mapsto A \mid \sigma \rangle \vdash a(x, b, p). P} \quad \text{I-ser}_2 \frac{\Delta \Downarrow \quad \langle \Delta; a \mapsto A; * \mapsto A, b, p \mid p^{\mathbb{O}} \rangle \vdash P}{\langle \Delta; a \mapsto A \mid \emptyset \rangle \vdash !a(x, b, p). P} \\
\text{I-con} \frac{\langle \Delta; * \mapsto A, b \mid q^{\mathbb{O}} \rangle \vdash P}{\langle \Delta; p \mapsto A \mid p^{\mathbb{I}}, q^{\mathbb{O}} \rangle \vdash p(x, b). P} \quad \text{RES-ser} \frac{\langle \Delta \mid \sigma \rangle \vdash P}{\langle \Delta^{-a} \mid \sigma \rangle \vdash \nu a P} \\
\text{RES-con}_1 \frac{\langle \Delta \mid \sigma_1, p^{\mathbb{O}}, p^{\mathbb{I}}, \sigma_2 \rangle \vdash P}{\langle \Delta^{-p} \mid \sigma_1, \sigma_2 \rangle \vdash \nu p P} \quad \text{RES-con}_2 \frac{\langle \Delta \mid \sigma \rangle \vdash P \quad p \notin \sigma}{\langle \Delta^{-p} \mid \sigma \rangle \vdash \nu p P} \\
\text{PAR} \frac{\langle \Delta_1 \mid \sigma_1 \rangle \vdash P_1 \quad \langle \Delta_2 \mid \sigma_2 \rangle \vdash P_2 \quad (\Delta_1, \Delta_2) \in \text{split}(\Delta) \quad \sigma \in \text{inter}(\sigma_1, \sigma_2)}{\langle \Delta \mid \sigma \rangle \vdash P_1 \mid P_2} \quad \text{NIL} \frac{\Delta \Downarrow}{\langle \Delta \mid \emptyset \rangle \vdash \mathbf{0}}
\end{array}$$

Figure 3: Typing rules with visibility and stack

Definition 26 (Interleaving). We write $\sigma_1 \in \text{inter}(\sigma_2; \sigma_3)$ if (i) σ_1 is a stack, and (ii) σ is an interleaving of σ_2 and σ_3 as by the following inductive rules:

1. $\emptyset \in \text{inter}(\emptyset; \emptyset)$
2. $p : \circ, \sigma_1 \in \text{inter}(\sigma_2; p : \circ, \sigma_3)$ if $\sigma_1 \in \text{inter}(\sigma_2; \sigma_3)$
3. $p : \circ, \sigma_1 \in \text{inter}(p : \circ, \sigma_2; \sigma_3)$ if $\sigma_1 \in \text{inter}(\sigma_2; \sigma_3)$
4. the same as (2) and (3) with $p : \mathbf{i}$ instead of $p : \circ$

If a name appears both in σ_2 and in σ_3 with the same tag, then $\text{inter}(\sigma_2; \sigma_3)$ is empty. Figure 3 gives the typing rules. The rules combine the handling of visibility and the update of the stack. For instance, in rule **0-SER**, newly created names b and p are added to the visibility when typing the continuation P , and $p^{\mathbb{I}}$ is added on top of the stack, meaning that P shall perform an input at p before interacting on other continuation names. The output capability on p is transmitted along the output at a ; correspondingly, in rules **I-SER₁** and **I-SER₂**, $p^{\mathbb{O}}$ is added to σ . In contrast with [3], we do not impose $\sigma = \emptyset$ in rule **I-SER₁**, as we do not enforce input receptiveness on continuation names. Performing inputs and outputs at continuation names has the effect of consuming the corresponding particle on top of the stack. The stack is empty when typing a replicated input, so to avoid duplication of continuation names. In rules **I-SER₂** and **NIL** we need to check that Δ is duplicable to make sure that all the input on continuations are defined. Rules **SUCC**, **0-F0**, **I-F0**, **RES-F0**, **SUM** and **IF** are like in Figure 1, and are omitted.

Theorem 27 (Subject reduction). *Suppose $\langle \Delta \mid \sigma \rangle \vdash P$ and $P \xrightarrow{\mu} P'$. We have:*

- if $\mu = \tau$, then either $\sigma = p^o, p^I, \sigma''$ with $p \in \text{Ncon} \setminus \text{fn}(P')$, in which case $\langle \Delta^{-p} \mid \sigma'' \rangle \vdash P'$,
or $\langle \Delta \mid \sigma \rangle \vdash P'$;
- if $\mu = \bar{a}\langle v \rangle(b, p)$ then $a \in \Delta(*)$ and $\langle \Delta^{-*}; b \mapsto \Delta(*), b; p \mapsto \Delta(*), b \mid p^I, \sigma \rangle \vdash P'$;
- if $\mu = \bar{p}\langle v \rangle(b)$ then $\sigma = p^o, \sigma''$, $p \in \Delta(*)$ and $\langle \Delta^{-*}; b \mapsto \Delta(*), b \mid \sigma'' \rangle \vdash P'$;
- if $\mu = a\langle v \rangle(b, p)$ and $\Delta \not\ll *$ then $\langle \Delta; * \mapsto \Delta(a), b, p \mid p^o, \sigma \rangle \vdash P'$;
- if $\mu = p\langle v \rangle(b)$ and $\Delta \not\ll *$ then $\sigma = p^I, q^o, \sigma''$ and $\langle \Delta^{-p}; * \mapsto \Delta(p), b \mid q^o, \sigma'' \rangle \vdash P'$;
- if $\mu = h\langle v \rangle$ or $(\mu = \bar{h}\langle v \rangle$ and $\Delta \not\ll *)$ then $\langle \Delta \mid \sigma \rangle \vdash P'$.

As previously, we have to define compatibility of typing environments so to be able to use the contextual forms of behavioural equivalence from Section 2.

Definition 28. $\langle \Delta_1 \mid \sigma_1 \rangle$ and $\langle \Delta_2 \mid \sigma_2 \rangle$ are compatible if $(\Delta_1 \cap \Delta_2) \ddagger$ and $\text{inter}(\sigma_1, \sigma_2) \neq \emptyset$.

Example 29. Consider the process

$$P \stackrel{\text{def}}{=} \nu h (\bar{h}\langle 5 \rangle \mid \bar{a}(b, q). (!b(p). h(x). (\bar{h}\langle x+1 \rangle. \mathbf{0} \mid \bar{p}. \mathbf{0}) \mid q. h(x). \bar{h}\langle x \rangle. \bar{c}(d, r). r. h(y). \bar{h}\langle y \rangle. \text{if } x = y \text{ then } \bar{s}. \mathbf{0} \text{ else } \mathbf{0}_*))$$

Here h is used as an integer reference, initially set to 5. The process calls a service a , giving access to a name b that can be used to increment the reference. When the call at a returns (at q), the reference is read twice and, in between, a call at c is made. Without visibility, the two values read for h might be different: as b has been exported, the environment could freely use b and increment the reference in between the two reads. Visibility ensures us that, when the initial call at a is returned, name b is ‘forgotten’ by the environment, which cannot therefore use it anymore. Thus, if Q is the process defined like P but with the last line replaced with $q. \bar{c}(d, r). r. \bar{s}. \mathbf{0}$, then we have $P \cong Q$. We refer to Appendix D.1 for an ML-like presentation of the example above, as well as for a longer example, also showing how re-entrant calls may sometimes make visibility boundaries hard to predict.

5.2 Labelled semantics

To represent the point of view of the observer, we use typing environments of the form $[\Gamma \mid \kappa]$, where Γ is a visibility function and κ is a *stack*. Figure 4 defines the typed transitions, written $[\Gamma \mid \kappa]P \xrightarrow{\mu} [\Gamma' \mid \kappa']P'$. Rules **TrTau**, **Tr0-fo** and **TrI-fo** are similar to their counterpart in Figure 2, and are omitted.

A trace $\mu_1 \dots \mu_n$ is a *well-bracketed trace* of P_0 under $[\Gamma_0 \mid \kappa_0]$ if there are P_i, Γ_i, κ_i ($1 \leq i \leq n$) with $[\Gamma_i \mid \kappa_i]P_i \xrightarrow{\mu_i} [\Gamma_{i+1} \mid \kappa_{i+1}]P_{i+1}$ for all $0 \leq i < n$.

$$\begin{array}{c}
\text{Tr0-ser} \frac{P \xrightarrow{\bar{a}\langle v \rangle(b,p)} P' \quad a \in \text{dom}(\Gamma) \quad \Gamma \not\vdash *}{[\Gamma \mid \kappa]P \xrightarrow{\bar{a}\langle v \rangle(b,p)} [\Gamma; * \mapsto \Gamma(a), b, p \mid p^0, \kappa]P'} \\
\text{Tr0-con} \frac{P \xrightarrow{\bar{p}\langle v \rangle(b)} P' \quad p \in \text{dom}(\Gamma) \quad \Gamma \not\vdash *}{[\Gamma \mid p^I, q^0, \kappa]P \xrightarrow{\bar{p}\langle v \rangle(b)} [\Gamma^{-p}; * \mapsto \Gamma(p), b \mid q^0, \kappa]P'} \\
\text{TrI-ser} \frac{P \xrightarrow{a\langle v \rangle(b,p)} P' \quad a \in \Gamma(*)}{[\Gamma \mid \kappa]P \xrightarrow{a\langle v \rangle(b,p)} [\Gamma^{-*}; p \mapsto \Gamma(*), b; b \mapsto \Gamma(*), b \mid p^I, \kappa]P'} \\
\text{TrI-con} \frac{P \xrightarrow{p\langle v \rangle(b)} P' \quad p \in \Gamma(*)}{[\Gamma \mid p^0, \kappa]P \xrightarrow{p\langle v \rangle(b)} [\Gamma^{-*}; b \mapsto \Gamma(*), b \mid \kappa]P'}
\end{array}$$

Figure 4: Visible LTS with stack

Definition 30 (Typed trace equivalence). Two $\langle \Delta \mid \sigma \rangle$ -processes P and Q are *trace equivalent at* $[\Gamma \mid \kappa]$, written $P \simeq_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$, if $\langle \Delta \mid \sigma \rangle$ and $[\Gamma \mid \kappa]$ are compatible and P and Q have the same sets of well-bracketed traces under $[\Gamma \mid \kappa]$.

This typed version of trace equivalence enjoys properties, including substitutivity properties, similar to those of the trace equivalence in Section 4. We only report the characterisation result with respect to may testing.

Theorem 31 (Characterisation of may testing using trace equivalence). *Suppose $\langle \Delta \mid \sigma \rangle \vdash P, Q$; then $P \simeq_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$ if and only if $P \cong_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$.*

The analogous result for typed bisimilarity is in Appendix D, together with results about the shape of type-allowed traces akin to those in Section 4.2.

6 Conclusions and Future Work

In this paper we have studied the meaning of visibility, originating from game-semantics interpretations of typed λ -calculi and reflecting absence of higher-order store, in a calculus of pure name-passing such as the π -calculus. In contrast with λ -calculi, in the π -calculus there is no explicit notion of store, and no unique identities (e.g., the input end of a name may have multiple occurrences, arbitrarily structured). These aspects determine considerable differences in the semantic analysis of visibility. For instance, in the π -calculus visibility has to be enforced by means of a dedicated type system; and the visibility functions have to satisfy special properties such as transitivity. We have highlighted such differences, and investigated the behavioural effects of visibility. In particular, we have shown full abstract

characterisations of may testing and barbed equivalence, as form of trace equivalence and labelled bisimilarity.

There are several directions for future work that we would like to pursue. In game semantics, further aspects related to first-order store have been studied; for instance allowing first-order references as values (that can be passed around) [12, 8], or allowing to store first-order references (i.e., references with types such as `int ref ref` [16, 9]). We would like to investigate the effects of such distinctions on our types and proof techniques.

We have worked with πI , a dialect of π -calculus in which all names exchanged are newly created. Thus πI lacks the free-output construct, where the communication of existing names is permitted. As shown in the literature [4, 24, 7], πI is closer to game semantics than the standard π -calculus. We would like to see whether our results can be adapted so to accommodate also the free output construct. While free output can be modelled in πI [2], it allows a direct representation of idioms such as ‘delegate to someone else the task of answering’, or the expression of tail-call recursion.

We have studied visibility on top of type systems for sequentiality and well-bracketing. Allowing concurrency would probably require relaxing the definition of the visibility functions; we leave the details to future work.

References

- [1] Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniérou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Romyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral types in programming languages. *Found. Trends Program. Lang.*, 3(2-3):95–230, 2016. doi:10.1561/25000000031.
- [2] Michele Boreale. On the expressiveness of internal mobility in name-passing calculi. *Theor. Comput. Sci.*, 195(2):205–226, 1998. doi:10.1016/S0304-3975(97)00220-X.
- [3] Daniel Hirschhoff, Enguerrand Prebet, and Davide Sangiorgi. On sequentiality and well-bracketing in the π -calculus. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, June 2021. URL: <http://dx.doi.org/10.1109/LICS52264.2021.9470559>, doi:10.1109/lics52264.2021.9470559.
- [4] Kohei Honda. Processes and games. In Fabio Gadducci and Ugo Montanari, editors, *Fourth International Workshop on Rewriting logic and Its Applications, WRLA2002, Pisa, Italy, 19-21, 2002*, volume 71 of *Electronic Notes in Theoretical Computer Science*, pages 40–69. Elsevier, 2002. doi:10.1016/S1571-0661(05)82528-9.
- [5] Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In Chris Hankin, editor, *Programming Languages and Systems - ESOP’98, 7th European Symposium*

- on Programming, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings, volume 1381 of *Lecture Notes in Computer Science*, pages 122–138. Springer, 1998. URL: <https://doi.org/10.1007/BFb0053567>, doi:10.1007/BFb0053567.
- [6] Guilhem Jaber and Andrzej S. Murawski. Complete trace models of state and control. In *Programming Languages and Systems: 30th European Symposium on Programming, ESOP 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 – April 1, 2021, Proceedings*, page 348–374, Berlin, Heidelberg, 2021. Springer-Verlag. doi:10.1007/978-3-030-72019-3_13.
- [7] Guilhem Jaber and Davide Sangiorgi. Games, Mobile Processes, and functions. In *CSL 2022 - 30th EACSL Annual Conference on Computer Science Logic*, pages 1–35, Göttingen, Germany, February 2022. URL: <https://hal.science/hal-03407123>.
- [8] Guilhem Jaber and Nikos Tzevelekos. Trace semantics for polymorphic references. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 585–594. ACM, 2016. doi:10.1145/2933575.2934509.
- [9] Ohad Kammar, Paul B. Levy, Sean K. Moss, and Sam Staton. A monad for full ground reference cells. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, Reykjavik, Iceland, 2017. IEEE. doi:10.1109/LICS.2017.8005109.
- [10] Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. *ACM Trans. Program. Lang. Syst.*, 21(5):914–947, 1999. doi:10.1145/330249.330251.
- [11] Naoki Kobayashi and Davide Sangiorgi. A hybrid type system for lock-freedom of mobile processes. *ACM Trans. Program. Lang. Syst.*, 32(5):16:1–16:49, 2010. doi:10.1145/1745312.1745313.
- [12] James Laird. A fully abstract trace semantics for general references. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007, Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 667–679. Springer, 2007. doi:10.1007/978-3-540-73420-8_58.
- [13] 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.
- [14] Robin Milner. Functions as processes. *Mathematical Structures in Computer Science*, 2(2):119–141, June 1992. doi:10.1017/S0960129500001407.

- [15] Robin Milner. Functions as processes. *Math. Struct. Comput. Sci.*, 2(2):119–141, 1992. doi:10.1017/S0960129500001407.
- [16] Andrzej S. Murawski and Nikos Tzevelekos. Algorithmic games for full ground references. *Formal Methods in System Design*, 52(3):277–314, 2018. doi:10.1007/s10703-017-0292-9.
- [17] Luca Padovani. Deadlock and lock freedom in the linear π -calculus. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 72:1–72:10. ACM, 2014. doi:10.1145/2603088.2603116.
- [18] Benjamin C. Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. *Math. Struct. Comput. Sci.*, 6(5):409–453, 1996. URL: <https://doi.org/10.1017/s096012950007002x>, doi:10.1017/S096012950007002X.
- [19] Enguerrand Prebet. *Typed Behavioural Equivalences in the Pi-Calculus*. Theses, Ecole normale supérieure de lyon - ENS LYON ; Università degli studi (Bologna, Italie), September 2022. URL: <https://hal.science/tel-03920089>.
- [20] Davide Sangiorgi. The lazy lambda calculus in a concurrency scenario. *Inf. Comput.*, 111(1):120–153, 1994. URL: <https://doi.org/10.1006/inco.1994.1042>, doi:10.1006/INCO.1994.1042.
- [21] Davide Sangiorgi. π -calculus, internal mobility, and agent-passing calculi. *Theor. Comput. Sci.*, 167(1&2):235–274, 1996. doi:10.1016/0304-3975(96)00075-8.
- [22] Davide Sangiorgi. The name discipline of uniform receptiveness. *Theor. Comput. Sci.*, 221(1-2):457–493, 1999. doi:10.1016/S0304-3975(99)00040-7.
- [23] Nobuko Yoshida, Martin Berger, and Kohei Honda. Strong normalisation in the π -calculus. *Inf. Comput.*, 191(2):145–202, 2004. URL: <https://doi.org/10.1016/j.ic.2003.08.004>, doi:10.1016/J.IC.2003.08.004.
- [24] Nobuko Yoshida, Simon Castellan, and Léo Stefanescu. Game semantics: Easy as π . *CoRR*, abs/2011.05248, 2020. URL: <https://arxiv.org/abs/2011.05248>, arXiv:2011.05248.

A Section 2: Definition of the Operational Semantics

Structural congruence, written \equiv , is the smallest relation that is an equivalence relation, contains α -conversion, is preserved by all operators of the calculus, and satisfies the following

axioms:

$$\begin{array}{c}
\frac{}{P \mid Q \equiv Q \mid P} \quad \frac{}{P \mid (Q \mid R) \equiv (P \mid Q) \mid R} \quad \frac{}{P \mid 0 \equiv P} \quad \frac{}{P + Q \equiv Q + P} \\
\\
\frac{}{P + (Q + R) \equiv (P + Q) + R} \quad \frac{}{P + 0 \equiv P} \quad \frac{a \notin \text{fn}(Q)}{(\nu a P) \mid Q \equiv \nu a(P \mid Q)} \\
\\
\frac{}{\nu a \nu b P \equiv \nu b \nu a P} \quad \frac{}{! \alpha. P \equiv ! \alpha. P \mid \alpha. P}
\end{array}$$

Similar versions of axioms for restrictions of the form νh instead of νa have been omitted.

To define the labelled transition system (LTS), we introduce *actions*, defined by the following grammar:

$$\mu ::= \tau \mid a\langle v \rangle(b) \mid \bar{a}\langle v \rangle(b) \mid h\langle v \rangle \mid \bar{h}\langle v \rangle \mid \omega$$

Actions are *early* for first-order values, and *late* for the transmission of HO names. We write $\text{bn}(\mu)$ (resp. $\text{fn}(\mu)$) for the bound (resp. free) names of action μ , defined by saying that b is a bound name in actions $a\langle v \rangle(b)$ and $\bar{a}\langle v \rangle(b)$, and the other names in actions are free.

We suppose that any closed first-order expression e can be evaluated to some value v . We write $e \Downarrow v$ for the corresponding evaluation relation. We furthermore suppose that equality between values can be tested.

The LTS is defined, on closed processes, by the rules in Figure 5.

We write $P \downarrow_\omega$ if $P \xrightarrow{\omega}$, i.e., if P contains an unguarded occurrence of an ω prefix at toplevel. We write $P \not\downarrow_\omega$ when this is not the case.

B Additional Material for Section 3

B.1 Properties of the type system

Lemma 32. *Suppose $\Delta \vdash P$ and $a \notin \text{fn}(P)$. Then also $\Delta^{-a} \vdash P$.*

Lemma 33 (Invariance of typing w.r.t. structural congruence). *If $\langle \Delta \mid \sigma \rangle \vdash P$ and $P \equiv P'$, then also $\langle \Delta \mid \sigma \rangle \vdash P'$.*

B.2 Properties of trace equivalence

Lemma 34 (Shrinking). *Suppose $\Gamma' \subseteq \Gamma$, Δ and Γ' compatible and $P \simeq_{\Delta}^{\Gamma} Q$. Then also $P \simeq_{\Delta}^{\Gamma'} Q$.*

Definition 35 (Projection). Let Γ be a visibility function and N a set of names. We define a visibility function as follows:

$$\text{proj}(N, \Gamma) : \begin{array}{ll} \text{HO} \cup \{*\} & \rightarrow \mathcal{P}(\text{HO}) \\ a \in \text{dom}(\Gamma) \cap N & \mapsto \Gamma(a) \cap N \end{array}$$

$$\begin{array}{c}
\frac{}{a(x, b). P \xrightarrow{a(v)(b)} P[v/x]} \quad \frac{e \Downarrow v}{\bar{a}\langle e \rangle(b). P \xrightarrow{\bar{a}(v)(b)} P} \quad \frac{}{h(x). P \xrightarrow{h(v)} P[v/x]} \\
\frac{e \Downarrow v}{\bar{h}\langle e \rangle. P \xrightarrow{\bar{h}(v)} P} \quad \frac{}{\omega. P \xrightarrow{\omega} \mathbf{0}} \quad \frac{\alpha. P \xrightarrow{\mu} Q}{!\alpha. P \xrightarrow{\mu} Q \mid !\alpha. P} \quad \frac{P \xrightarrow{\mu} Q}{\nu a P \xrightarrow{\mu} \nu a Q} \quad a \notin \text{fn}(\mu) \\
\frac{P \xrightarrow{\mu} Q}{\nu h P \xrightarrow{\mu} \nu h Q} \quad h \notin \text{fn}(\mu) \quad \frac{P \xrightarrow{\mu} P'}{P + Q \xrightarrow{\mu} P'} \quad \frac{P \xrightarrow{a(v)(b)} P' \quad Q \xrightarrow{\bar{a}(v)(b)} Q'}{P \mid Q \xrightarrow{\tau} \nu b(P' \mid Q')} \\
\frac{P \xrightarrow{h(v)} P' \quad Q \xrightarrow{\bar{h}(v)} Q'}{P \mid Q \xrightarrow{\tau} P' \mid Q'} \quad \frac{P \xrightarrow{\mu} P'}{P \mid Q \xrightarrow{\mu} P' \mid Q} \quad \text{fn}(Q) \cap \text{bn}(\mu) = \emptyset \\
\frac{e \Downarrow v \quad f \Downarrow v \quad P \xrightarrow{\mu} P'}{\text{if } e = f \text{ then } P \text{ else } Q \xrightarrow{\mu} P'} \quad \frac{e \Downarrow v \quad f \Downarrow v' \quad v \neq v' \quad Q \xrightarrow{\mu} Q'}{\text{if } e = f \text{ then } P \text{ else } Q \xrightarrow{\mu} Q'}
\end{array}$$

Figure 5: Labelled transition semantics.

Symmetric versions of the rules for parallel composition and for sum have been omitted.

Lemma 36 (Extension). *Consider $\Delta, \Gamma, \Gamma', P, Q$ such that $\text{proj}(\text{fn}(\Gamma, P, Q), \Gamma') = \Gamma$ and Δ and Γ' are compatible. If $P \simeq_{\Delta}^{\Gamma} Q$ then $P \simeq_{\Delta}^{\Gamma'} Q$.*

In the statement above, $\text{fn}(\Gamma, P, Q)$ stands for the union of all names mentioned in Γ , together with the names occurring free in P and Q .

B.3 Pruning Processes using Types

Consider a visibility function Δ and $S \subseteq \mathbf{H0}$. We write $\alpha \wedge S$ when α is an input of the form $a(x, b)$ and either $a \in S$ or $(a \in \text{dom}(\Delta)$ and $\Delta(a) \cap S \neq \emptyset$).

We define inductively the pruning function on processes that removes all instances of the names in S as follows:

$$\Pr_S^{\Delta}(\alpha. P) = \begin{cases} \mathbf{0} & \text{if } \alpha \wedge S \\ \alpha. \Pr_S^{\Delta}(P) & \text{otherwise} \end{cases} \quad \Pr_S^{\Delta}(!\alpha. P) = \begin{cases} \mathbf{0} & \text{if } \alpha \wedge S \\ !\alpha. \Pr_S^{\Delta}(P) & \text{otherwise} \end{cases}$$

$$\Pr_S^{\Delta}(\bar{\alpha}. P) = \bar{\alpha}. \Pr_S^{\Delta}(P) \quad \Pr_S^{\Delta}(\nu a P) = \nu a \Pr_S^{\Delta}(P) \quad \Pr_S^{\Delta}(\nu h P) = \nu h \Pr_S^{\Delta}(P)$$

$$\Pr_S^{\Delta}(P \mid Q) = \Pr_S^{\Delta}(P) \mid \Pr_S^{\Delta}(Q) \quad \Pr_S^{\Delta}(P + Q) = \Pr_S^{\Delta}(P) + \Pr_S^{\Delta}(Q) \quad \Pr_S^{\Delta}(\mathbf{0}) = \mathbf{0}$$

$$\Pr_S^{\Delta}(\omega) = \omega \quad \Pr_S^{\Delta}(\text{if } e = f \text{ then } P \text{ else } Q) = \text{if } e = f \text{ then } \Pr_S^{\Delta}(P) \text{ else } \Pr_S^{\Delta}(Q)$$

Lemma 37 (Pruning). *Let Γ, Δ, P such $\Delta \vdash P$ and Γ and Δ are compatible and $\Delta \cup \Gamma$ is transitive. Then for all $S \subseteq \mathbf{HO}$ such that $S \cap (\Gamma \cup \Delta(*)) = \emptyset$, $\Delta^{-S} \vdash \text{Pr}_S^\Delta(P)$ and $P \simeq_\Delta^\Gamma \text{Pr}_S^\Delta(P)$.*

We actually remove more than the names in S . Indeed, to preserve sequentiality, we cannot simply remove an output; we actually remove all inputs that may cause an output on the names in S .

The idea of the proof is to show that $S \cap (\Gamma \cup \Delta(*)) = \emptyset$ is an invariant throughout the trace; for this we use repeatedly transitivity of $\Gamma \cup \Delta$.

Proof. Let $\mu_1 \dots \mu_n$ a trace of P under Γ , let P_i, Γ_i be the processes and visibility functions given that trace and Δ_i the visibility functions given by subject reduction. We show by induction on i that there exists Γ'_i, Δ'_i such that $\Gamma'_i \supseteq \Gamma_i$ and $\Delta'_i \supseteq \Delta_i$ and $S \cap (\Gamma'_i \cup \Delta'_i)(*) = \emptyset$ and that Δ'_i and Γ'_i are compatible, have a transitive union and coincide on the intersection of their domains. The case $i = 0$ is immediate, let $0 \leq i < n$, we assume the induction hypothesis on i and show it for $i + 1$ by cases on μ_i :

- if $\mu_i = \tau$, or $\mu_i = h\langle v \rangle$ or $\mu_i = \bar{h}\langle v \rangle$ then $\Gamma_{i+1} = \Gamma_i$ and $\Delta_{i+1} = \Delta_i$, these cases are immediate by setting $\Gamma'_{i+1} = \Gamma'_i$ and $\Delta'_{i+1} = \Delta'_i$.
- if $\mu_i = \bar{a}\langle v \rangle(b)$ then $\Gamma_{i+1} = \Gamma_i; * \mapsto \Gamma_i(a), b$ and $\Delta_{i+1} = \Delta_i^{-*}; b \mapsto \Delta_i(*), b$. Let $\Gamma'_{i+1} = \Gamma'_i; * \mapsto \Delta'_i(*), b \supseteq \Gamma_{i+1}$ because $a \in \Delta_i(*)$ therefore $\Gamma_i(a) \subseteq \Delta_i(*)$. And let $\Delta'_{i+1} = \Delta_i'^{-*}; b \mapsto \Delta'_i(*), b$.
 Γ'_{i+1} and Δ'_{i+1} are compatible, their union is transitive and they coincide on the intersection of their domains.
 $(\Gamma'_{i+1} \cup \Delta'_{i+1})(*) = (\Gamma'_i \cup \Delta'_i)(*), b$ where b is fresh therefore $b \notin S$ and by induction hypothesis $S \cap (\Gamma'_{i+1} \cup \Delta'_{i+1})(*) = \emptyset$.
- if $\mu_i = a\langle v \rangle(b)$ the case is similar to the previous one, swapping Γ and Δ .

This ensures us that none of the μ_i are inputs or outputs on higher-order names in S . Moreover if $a \in \text{dom}(\Delta)$ and $\Delta(a) \cap S \neq \emptyset$ then for all i , $a \notin (\Gamma_i \cup \Delta_i)(*)$ otherwise we would have $\emptyset \neq \Delta(a) \cap S = \Delta_i(a) \cap S \subseteq (\Gamma_i \cup \Delta_i)(*) \cap S = \emptyset$.

Therefore the trace $\mu_1 \dots \mu_n$ is also a trace of $\text{Pr}_S^\Delta(P)$ under Γ . □

C Section 4: Labelled Bisimilarity

A *typed relation* is a set of triples (Γ, P, Q) intended to represent the visibility of the observer for processes P and Q in the bisimulation game.

Definition 38 (Typed Bisimulation). A typed relation \mathcal{R} is a *typed (weak) bisimulation* if whenever $(\Gamma, P, Q) \in \mathcal{R}$:

1. if $[\Gamma] P \xrightarrow{\mu} [\Gamma'] P'$ then there is Q' such that $\Gamma Q \xrightarrow{\hat{\mu}} \Gamma' Q'$ and $(\Gamma', P', Q') \in \mathcal{R}$;

2. the converse, on the actions from Q .

Typed bisimilarity, written \approx , is the largest typed bisimulation. For two Δ -processes P, Q we write $P \approx_{\Delta}^{\Gamma} Q$ when $(\Gamma, P, Q) \in \approx$ and Δ and Γ are compatible. In this case, we say that P and Q are *bisimilar under Γ* .

Example 39. As for barbed equivalence the visibility of opponent also plays an important role in determining whether two process are bisimilar under some Γ . Consider Δ and Γ compatible, and a Δ -process of the form $Q = P \mid a.\bar{b}.R$ with $b \notin \text{fn}(P)$ and $b \in \text{dom}(\Gamma)$:

- if $a \in \text{fn}(P)$, we may have, for example, $P = \bar{a}.R'$, in which case $Q \xrightarrow{\tau} R' \mid \bar{b}.R$; such derivative is not bisimilar with P under Γ , as only $R' \mid \bar{b}.R$ can make an output on b ;
- if $a \in \Gamma(*)$, then we have $Q \xrightarrow{a} P \mid \bar{b}.R$, which is not bisimilar to any P' such that $P \xrightarrow{a} P'$, for the same reason as in the previous case;
- however, if both $a \notin \text{fn}(P)$ and $a \notin \Gamma(*)$, then Q cannot make an input on a , and therefore we have $Q \approx_{\Delta}^{\Gamma} P$.

The analogues of Lemmas 16-19 and 37 also hold for typed bisimilarity.

Lemma 40 (Restriction). *If $P \approx_{\Delta}^{\Gamma} Q$, then also $\nu a P \approx_{\Delta-a}^{\Gamma-a} \nu a Q$.*

Lemma 41 (Input). *If $P \approx_{\Delta}^{\Gamma} Q$ and $\Gamma \not\# *$, then also $a(x, b).P \approx_{\Delta-x, b}^{(\Gamma; * \mapsto \Gamma(b))^{-b}} a(x, b).Q$.*

Lemma 42 (Output). *If $P \approx_{\Delta}^{\Gamma} Q$ and $\Gamma \not\# *$, then also $\bar{a}\langle e \rangle(b).P \approx_{(\Delta; * \mapsto \Delta(b))^{-b}}^{\Gamma-*} \bar{a}\langle e \rangle(b).Q$.*

Lemma 43 (Parallel composition). *Suppose $P \approx_{\Delta}^{\Gamma} Q$ and $(\Gamma_1, \Gamma_2) \in \text{split}(\Gamma)$. Then $\Gamma_2 \vdash R$ implies $P \mid R \approx_{(\Delta \cup \Gamma_2)^+}^{\Gamma_1} Q \mid R$.*

Lemma 44 (Pruning). *Let Γ, Δ, P such $\Delta \vdash P$ and Γ and Δ are compatible. Then for all $S \subseteq \mathbf{H0}$ such that $S \cap (\Gamma \cup \Delta(*)) = \emptyset$, $\Delta^{-S} \vdash \text{Pr}_S^{\Delta}(P)$ and $P \approx_{\Delta}^{\Gamma} \text{Pr}_S^{\Delta}(P)$.*

Technical lemmas similar to those for trace equivalence also holds, such as the following ones.

Lemma 45 (Shrinking). *Suppose $\Gamma' \subseteq \Gamma$, Δ and Γ' compatible and $P \approx_{\Delta}^{\Gamma} Q$. Then also $P \approx_{\Delta}^{\Gamma'} Q$.*

Lemma 46 (extension). *Consider Γ, Γ', P, Q such that $\text{proj}(\text{fn}(\Gamma, P, Q), \Gamma') = \Gamma$ and Δ and Γ' are compatible. If $P \approx_{\Delta}^{\Gamma} Q$ then $P \approx_{\Delta}^{\Gamma'} Q$.*

We report the characterisation results. Soundness follows from the substitutivity properties of labelled bisimilarity. For completeness, as usual with characterisations of barbed equivalence, the proof uses induction, via an inductive stratification of bisimilarity, and requires image-finiteness of processes (that is, up-to α -conversion, the set of possible type-allowed weak transitions emanating from a process is finite). The image-finite condition is needed in the characterisation of bisimilarity via the inductive stratification. Also, and again in line with proofs about barbed relations, the proof exploits the presence of the sum operator and of success signals that can be emitted by a testing context whenever needed.

Theorem 47 (Soundness of bisimilarity for barbed equivalence). *Suppose $\Delta \vdash P, Q$ with $P \approx_{\Delta}^{\Gamma} Q$. Then also $P \cong_{\Delta}^{\Gamma} Q$.*

Theorem 48 (Completeness of bisimilarity for barbed equivalence). *If P and Q are image-finite processes, and if $P \cong_{\Delta}^{\Gamma} Q$, then $P \approx_{\Delta}^{\Gamma} Q$.*

D Additional Material for Section 5

D.1 About Example 29

The process discussed in Example 29 can be seen as the π -calculus counterpart of the following ML program, which might be more readable for a reader familiar with functional languages:

```

let h = ref 0 in
let b () = h := !h +1 in
begin
  a b;
  let x = !h in
  c x;
  let y = !h in
  if x=y then loop() else ()
end

```

Re-entrant call. The following example first shows an equality similar to that of Example 29; and then develops the example so to show how re-entrant calls may sometimes break ‘expected’ visibility boundaries.

We define two processes P_1 and P_2 , which access a reference h (implemented using a first-order name), and can call a server b . When calling b , the processes first pass the name `incr` of a server that increments the value stored in h . Once the call on b is answered the processes read the value in h and make another call to b passing the identity function. After that, the processes read again the value of h . Finally, P_1 outputs the first value read in h while P_2 outputs the second one.

$$\begin{aligned}
P_i \stackrel{\text{def}}{=} & \bar{b}(\text{incr}, q_1).q_1.h(x_1).(\bar{h}\langle x_1 \rangle \\
& | \bar{b}(\text{id}, q_2). \\
& q_2.h(x_2).(\bar{h}\langle x_2 \rangle \\
& | \bar{p}\langle x_i \rangle))
\end{aligned}$$

We first compare the processes

$$\nu h(\bar{h}\langle 0 \rangle | a(b, p).P_1 | Q_{\text{incr}} | Q_{\text{id}}) \quad \text{and} \quad \nu h(\bar{h}\langle 0 \rangle | a(b, p).P_2 | Q_{\text{incr}} | Q_{\text{id}}),$$

where Q_{incr} (resp. Q_{id}) is the encoding of function `incr` (resp. `id`).

Here, the only way to change the value stored in h is by calling `incr`, which is not visible by the observer after the input on q_1 . These two processes are bisimilar under any visibility.

Then we compare the processes

$$Q_1 \stackrel{\text{def}}{=} \nu h(\bar{h}\langle 0 \rangle \mid !a(b, p). P_1 \mid Q_{\text{incr}} \mid Q_{\text{id}}) \quad \text{and} \quad Q_2 \stackrel{\text{def}}{=} \nu h(\bar{h}\langle 0 \rangle \mid !a(b, p). P_2 \mid Q_{\text{incr}} \mid Q_{\text{id}}).$$

We are able to distinguish them using a re-entrant call :

$$R \stackrel{\text{def}}{=} \bar{a}(b, p). (b(\text{incr}, q). \bar{a}(d, r). (!d(-, s). \overline{\text{incr}}(t). t. \bar{s} \\ \mid r(n). \text{if } n = 1 \text{ then } \bar{q} \text{ else } \omega) \\ \mid p(m))$$

The idea is to call a a second time when `incr` is still in the observer's visibility; this means that the new server passed to a has access to `incr` during both calls. With the observer R , only Q_2 will be able to send a success signal. Therefore, for any visibility function Γ such that $a \in \Gamma(*)$, Q_1 and Q_2 are not bisimilar under visibility Γ .

The processes Q_1 and Q_2 can be seen as the π -calculus counterparts of the following ML program. (This example, with re-entrant calls, has been inspired by similar examples in game semantics of functional languages):

```
let h = ref 0 in
let incr () = h := !h + 1 in
let id x = x in
let a b =
  b incr;
  let x1 = !h in
  b id;
  let x2 = !h in
  xi
```

These programs can be distinguished by the following ML program that corresponds to the process R :

```
let b f =
  let d _ = f () in
  let n = a d in
  if n = 1 then () else  $\omega$ 
in
a b
```

The idea is the same as before: `d` calls `f` which will be `incr` in the first call to `b`; therefore at each call to `d` the reference is incremented. As a consequence, first program returns 1 while the second one returns 2.

D.2 Typed bisimilarity

Definition 49 (Typed Bisimulation). A typed relation \mathcal{R} is a *typed (weak) bisimulation* if whenever $([\Gamma \mid \kappa], P, Q) \in \mathcal{R}$:

1. if $[\Gamma \mid \kappa]P \xrightarrow{\mu} [\Gamma' \mid \kappa']P'$ then there is Q' such that $[\Gamma \mid \kappa]Q \xrightarrow{\hat{\mu}} [\Gamma' \mid \kappa']Q'$ and $([\Gamma' \mid \kappa'], P', Q') \in \mathcal{R}$;
2. the converse, on the actions from Q .

Typed bisimilarity, written \approx , is the largest typed bisimulation. For two $\langle \Delta \mid \sigma \rangle$ -processes P, Q we write $P \approx_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$ when $([\Gamma \mid \kappa], P, Q) \in \approx$ and $\langle \Delta \mid \sigma \rangle$ and $[\Gamma \mid \kappa]$ are compatible.

Theorem 50 (Soundness of bisimilarity for barbed equivalence). *Suppose $\langle \Delta \mid \sigma \rangle \vdash P, Q$ with $P \approx_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$. Then also $P \cong_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$.*

Theorem 51 (Completeness of bisimilarity for barbed equivalence). *If P and Q are image-finite processes, and if $P \cong_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$, then $P \approx_{\langle \Delta \mid \sigma \rangle}^{[\Gamma \mid \kappa]} Q$.*

D.3 Properties of traces: views and well-bracketing

Along the lines of the material presented in Section 4.2, we can reason about the traces produced by typable processes, based on the notion of justification. For this, we adapt Definition 22 to define views for traces. The notion of *respecting views* (Definition 23) is also adapted straightforwardly, which allows us to obtain the analogue of Lemma 24.

Lemma 52. *Consider $\Delta, \Gamma, \sigma, \kappa, P$ such that $\langle \Delta \mid \sigma \rangle \vdash P$ and $[\Gamma \mid \kappa]$ is compatible with $\langle \Delta \mid \sigma \rangle$. Let $A = (\Gamma \cup \Delta)(*)$, then every trace of P under $[\Gamma \mid \kappa]$ respects views with starting visibility A .*

As highlighted in [3], continuation names allow us to formulate a property of well-bracketing for traces, enforced by our type system. In contrast with [3], however, we do not impose input receptiveness on continuation names.

Definition 53 (Well-bracketing). An action of the form $a\langle v \rangle(b, p)$ or $\bar{a}\langle v \rangle(b, p)$ is called a *question*. An action of the form $p\langle v \rangle(b)$ or $\bar{p}\langle v \rangle(b)$ is called an *answer*.

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 $\mu_i \not\curvearrowright \mu_k$ and $\mu_k \not\curvearrowright \mu_j$ for all $i < k < j$, then $\mu_i \curvearrowright \mu_j$.

A stack σ is *clean* if each name has at most one occurrence, with any tag, in σ . A trace $\mu_0 \dots \mu_n$ is a *typed trace emanating from P_0* if there are $\Delta_1, \sigma_1, \dots, \Delta_n, \sigma_n, P_1, \dots, P_n$ such that for all $0 \leq j < n$ we are in one of the cases given by Theorem 27, which gives $\langle \Delta_{j+1} \mid \sigma_{j+1} \rangle$ when $P_j \xrightarrow{\mu_{j+1}} P_{j+1}$.

Lemma 54. *Suppose $\langle \Delta_0 \mid \sigma_0 \rangle \vdash P_0$, where σ_0 is clean. All typed traces emanating from P_0 are well-bracketed.*