Fair Termination of Binary Sessions

LUCA CICCONE, Università di Torino, Italy

LUCA PADOVANI, Università di Torino, Italy

A binary session is a private communication channel that connects two processes, each adhering to a protocol description called *session type*. In this work, we study the first type system that ensures the *fair termination* of binary sessions. A session fairly terminates if all of the infinite executions admitted by its protocol are deemed “unrealistic” because they violate certain *fairness assumptions*. Fair termination entails the eventual completion of all pending input/output actions, including those that depend on the completion of an unbounded number of other actions in possibly different sessions. This form of *lock freedom* allows us to address a large family of natural communication patterns that fall outside the scope of existing type systems. Our type system is also the first to adopt *fair subtyping*, a liveness-preserving refinement of the standard subtyping relation for session types that so far has only been studied theoretically. Fair subtyping is surprisingly subtle not only to characterize concisely but also to use appropriately, to the point that the type system must carefully account for all usages of fair subtyping to avoid compromising its liveness-preserving properties.

CCS Concepts: • **Theory of computation** → **Process calculi**; **Type structures**; **Program analysis**.

Additional Key Words and Phrases: session types, fair termination, fair subtyping, deadlock freedom

ACM Reference Format:

Luca Ciccone and Luca Padovani. 2022. Fair Termination of Binary Sessions. *Proc. ACM Program. Lang.* 6, POPL, Article 5 (January 2022), 30 pages. <https://doi.org/10.1145/3498666>

1 INTRODUCTION

Session type systems [Honda 1993; Honda et al. 1998; Hüttel et al. 2016] are an established formalism for the static analysis of communicating processes: a *binary session* is a private communication channel that connects two processes, each using one *endpoint* of the session; a *session type* is a type-level description of the sequences of input/output actions performed by a process with respect to a session endpoint. By making sure that the session types associated with the endpoints of a session complement each other and that processes do behave according to these types, it is possible to design type systems that enforce fundamental correctness properties such as communication safety, protocol fidelity, race and deadlock freedom. These are all instances of *safety properties*, guaranteeing that “nothing bad ever happens”, but in general one is also interested in *liveness properties* guaranteeing that “something good eventually happens” [Owicki and Lamport 1982].

To illustrate a few examples of liveness properties, consider the process

$$A\langle x \rangle \mid B\langle x, y \rangle \mid C\langle y \rangle \quad \text{where} \quad \begin{aligned} B(x, y) &\triangleq x? \{ \text{add} : B\langle x, y \rangle, \text{pay} : \dots y! \text{ship} \dots \} \\ C(y) &\triangleq y? \text{ship} \dots \end{aligned} \quad (1)$$

that models an acquirer A purchasing items from a business B which interacts with a carrier C . The acquirer and the business are connected by a session x whereas the business and the carrier are connected by a session y . The process is intentionally incomplete, but we see that the business can

Authors’ addresses: Luca Ciccone, Dipartimento di Informatica, Università di Torino, Torino, Italy, luca.ciccone@unito.it; Luca Padovani, Dipartimento di Informatica, Università di Torino, Torino, Italy, luca.padovani@unito.it.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/1-ART5

<https://doi.org/10.1145/3498666>

receive from the acquirer an arbitrary number of *add* messages (each message modeling the fact that the acquirer has added an item to its shopping cart) or a single *pay* message. Only then the business sends the *ship* message and the carrier can make progress. Examples of liveness properties are (L1) *the acquirer eventually sends pay to the business*, (L2) *the business eventually sends ship to the carrier*, and also (L3) *the sessions x and y eventually terminate*. Note that some of these properties are stronger than others. For example, if we know that (L3) holds, then both (L1) and (L2) hold too.

A type system ensuring properties like (L1)–(L3) for the process in Eq. (1) is missing. There exist (session) type systems ensuring somewhat related properties such as *deadlock freedom* [Caires et al. 2016; Dardha and Gay 2018; Padovani 2014; Wadler 2014], *lock freedom* [Kobayashi 2002; Kobayashi and Sangiorgi 2010; Padovani 2014; Scalas and Yoshida 2019] and *strong normalization* [Lindley and Morris 2016; Yoshida et al. 2004]. However, deadlock freedom is in general too weak to guarantee liveness properties such as (L1)–(L3). For example, if the acquirer keeps adding items to the cart but is never willing to pay the business, the process in Eq. (1) is deadlock-free even if the carrier cannot make any progress. At the same time, strong normalization implies (L1)–(L3) but is in general too strong, in the sense that there exist interesting processes that satisfy these properties but that are *not* strongly normalizing. For example, if the acquirer may buy arbitrarily many items, then the process in Eq. (1) admits, at least in principle, an infinite computation in which the acquirer keeps adding items to the cart even if it is willing to pay the business every now and then. Lock freedom is a liveness property meaning that every pending communication, such as the input $y?$ *ship*, can eventually be completed. Hence, lock freedom is strong enough to entail e.g. (L2). However, currently available type systems ensuring lock freedom have an important limitation: they can only handle those processes in which the completion of pending actions on a channel is attainable *regardless of the content of messages exchanged in other channels*. This is not the case for the action $y!$ *ship* on session y , which is performed only provided that a *pay* message is exchanged in session x . Note that B , far from being a contrived corner case, models a simple “while” loop that talks to A for an *arbitrarily long but supposedly finite amount of time* before turning the attention to C . In summary, Eq. (1) is representative of a large family of processes for which no available type system is able to provide strong liveness guarantees such as (L1)–(L3).

The type system we propose in this work ensures the fair termination of binary sessions. In general, fair termination means that a program terminates provided that some *fairness assumptions* are made on its executions [Francez 1986; Grumberg et al. 1984]. In our context, fair termination guarantees the eventual completion of every session under such assumptions. Fair termination is stronger than (dead)lock freedom but weaker than strong normalization. Unlike deadlock freedom, fair termination guarantees that every pending action may be completed. Unlike lock freedom, fair termination guarantees that every session may come to an end. Unlike strong normalization, fair termination does not rule out the existence of infinite executions, which are deemed “unrealistic” because they violate the fairness assumptions being made. There are two reasons why we focus on fair termination instead of lock freedom. First, the word “session” embodies the idea of an activity that lasts for a *finite period of time*, so (fair) session termination may be regarded as a desirable if not defining property of communicating sessions in the first place. Second, lock freedom does not scale well to multiple (chained, nested or interleaved) sessions. For instance, the stubborn acquirer that sends *add* messages forever results in a lock-free session x , but the action $y!$ *ship* in the business that is meant to “unlock” the complementary action $y?$ *ship* in the carrier can be performed only if the acquirer eventually sends *pay*. So, knowing that the session x is lock free does not help us to reason on the lock freedom of the session y . On the contrary, knowing that x fairly terminates is enough to deduce that *pay* can be sent, hence that y fairly terminates as well.

Among the several fairness notions that have been considered in the literature [Francez 1986; Kwiatkowska 1989; van Glabbeek and Höfner 2019] we assume *strong fairness*, namely we assume

that a process that has infinitely many opportunities of making some choice will make that choice infinitely often. In the specific case of (1), this translates to the assumption that an acquirer periodically faced with the opportunity of paying the business eventually pays the business. The motivation for choosing strong fairness is that we need an assumption on the individual behavior of sequential processes and on the messages they choose to send, whereas weaker assumptions like *e.g.* justness [van Glabbeek 2019; van Glabbeek and Höfner 2019] only concern the way parallel processes may independently progress. The fairness assumption we make is an assumption in a literal sense; it cannot be guaranteed by the type system or by a *scheduler* [Apt et al. 1987; Francez 1986] because it concerns the internal behavior of the processes that partake in a session.

It should be noted that the mere *assumption* of strong fairness does not turn an ordinary session type system into one that ensures fair session termination because the correspondence imposed by the type system between the structure of processes and that of the protocols they implement is generally (often necessarily) a loose one. Indeed, processes may be “more accommodating” than the protocols they implement by handling more messages than those mentioned in the protocols. For example, the business in (1) could handle a *search* message in addition to *add* and *pay*, even if the session type associated with x does not mention *search*. At the same time, processes may also be “less demanding” than the protocols they implement by sending fewer messages than those allowed by the protocols. For example, the acquirer in (1) could always purchase an odd number of items, or at least n items, or no more than n items, even if the session type associated with x allows sending an arbitrary number of *add* messages. These mismatches between processes and protocols are usually reconciled by a *subtyping relation* for session types [Bernardi and Hennessy 2016; Gay and Hole 2005]. The problem is that this subtyping relation is *too coarse* because it has been conceived to preserve the *safety* properties of sessions but not termination, which is a *liveness* property: if session types are not sufficiently precise descriptions of the actual behavior of processes, a session that appears to be fairly terminating at the level of types may not terminate at all at the level of processes. To solve this problem we adopt *fair subtyping* [Bravetti et al. 2021; Padovani 2013, 2016], a *liveness-preserving* refinement of the subtyping relation defined by Gay and Hole [2005].

As it turns out, the strong fairness assumption and the adoption of fair subtyping are still insufficient to guarantee fair session termination. For example, a process could indefinitely delay the termination of a session if it is allowed to chain or nest an infinite number of other sessions, even if all the created sessions (fairly) terminate. An even subtler issue is that *fair subtyping can be easily abused*, in the sense that using it “infinitely often” in the typing derivation of a recursive process may compromise its liveness-preserving feature. To overcome these problems, the type system has to account for all the creations of new sessions and all the usages of fair subtyping, making sure that the overall effort required to terminate all open sessions remains finite.

Summary of contributions. We present a session type system that ensures the *fair termination* of binary sessions in a calculus that supports *general recursion*, *session interleaving*, *session delegation* and *dynamic session creation*. Fair termination implies that all pending communications, including those that are blocked by an unbounded number of other communications and that depend on the exchange of a particular message in possibly different sessions, can be completed in finite time. To the best of our knowledge, this is the first type system capable of ensuring this form of lock freedom for a family of processes large enough to include (1). We also solve the long-standing problem of devising a type system based on *fair subtyping*. We design the type system so that the usage of fair subtyping in a typing derivation is safe, uncovering a dangerous interaction between fair subtyping and delegation whereby a single usage of fair subtyping for higher-order session types may have the same overall effect of infinitely many usages. We show how to avoid this issue by reconsidering some established properties of (fair) subtyping for higher-order session types.

Structure of the paper. Section 2 provides a quick introduction to *generalized inference systems*, the formalism we use to define fair subtyping and the typing rules. Section 3 provides all the notions on session types that are necessary in this work, including fair subtyping and fair termination. Section 4 describes the process calculus and Section 5 motivates the key properties enforced by the type system through a series of examples. Section 6 formalizes the type system and states its soundness. In Section 7 we show the dangerous interaction between fair subtyping and higher-order session types. Section 8 provides a more detailed comparison with related work and Section 9 concludes. Proofs and further definitions and results can be found in the Appendix, which is published in the supplemental material section of the ACM Digital Library page for this paper.

2 GENERALIZED INFERENCE SYSTEMS IN A NUTSHELL

Inference systems [Aczel 1977] are ubiquitous in the definition of predicates, relations and typing rules. An *inference system* $\mathcal{I}$ over a *universe* $\mathcal{U}$ of *judgments* is a set of *rules* $\langle pr, j \rangle$ where $pr \subseteq \mathcal{U}$ is a set of *premises* and $j \in \mathcal{U}$ is the *conclusion*. A *derivation tree* of $\mathcal{I}$ is a tree such that each node is labeled with the conclusion of a rule in $\mathcal{I}$ and its children are labeled with the premises of the rule. We say that a judgment j is *derivable* in $\mathcal{I}$ if there exists a derivation tree with root j . An inference system can have different interpretations depending on the set of derivation trees that one considers. The *inductive interpretation* of an inference system is the set of judgments that are derivable with well-founded derivation trees, those having a finite depth. The *coinductive interpretation* of an inference system is the set of judgments that are derivable with arbitrary (finite- or infinite-depth) derivation trees. It is known that these two interpretations respectively coincide with the least and the greatest fixed points of the inference operator $\Phi_{\mathcal{I}}$ associated with the inference system $\mathcal{I}$:

$$\Phi_{\mathcal{I}}(X) \stackrel{\text{def}}{=} \{j \in \mathcal{U} \mid \exists pr \subseteq X : \langle pr, j \rangle \in \mathcal{I}\} \quad \text{for all } X \subseteq \mathcal{U}$$

In some cases, the desired set of derivable judgments is an intermediate fixed point of the inference operator other than the least/greatest one. Generalized inference systems [Ancona et al. 2017; Dagnino 2019] allow for the characterization of (some) intermediate fixed points. Specifically, a *generalized inference system* is a pair $(\mathcal{I}, \mathcal{I}_{\text{co}})$ of inference systems whose interpretation is the set of judgments having an arbitrary (finite- or infinite-depth) derivation tree using the rules in $\mathcal{I}$ but such that all the judgments in this derivation tree also have a finite-depth derivation tree using the rules in $\mathcal{I} \cup \mathcal{I}_{\text{co}}$. It can be shown that this interpretation coincides with the greatest fixed point of $\Phi_{\mathcal{I}}$ that is included in the least fixed point of $\Phi_{\mathcal{I} \cup \mathcal{I}_{\text{co}}}$. The elements of $\mathcal{I}_{\text{co}}$ are called *corules*.

Hereafter, we write (co)rules following standard conventions: we use *meta-variables* for specifying families of (co)rules in a compact way and we draw a horizontal line to separate the premises pr from the conclusion j of a rule $\langle pr, j \rangle$. We double the line to distinguish the corules.

Example 2.1. The generalized inference system below defines a predicate $\text{maximum}(l, x)$ asserting that x is the greatest element of a *possibly infinite* list l :

$$\frac{}{\text{maximum}(x :: [], x)} \quad \frac{\text{maximum}(l, y)}{\text{maximum}(x :: l, \max\{x, y\})} \quad \frac{}{\text{maximum}(x :: l, x)}$$

The axiom on the left states that the greatest element of a list $x :: []$ that contains only x is just x , whereas the rule in the middle states that the greatest element of a list $x :: l$ is the maximum among the head x and the greatest element of the tail l . The problem of these “plain” rules is that their inductive interpretation is sound but not complete (no infinite list has a greatest element), whereas their coinductive interpretation is complete but not sound (it is possible to derive the judgment $\text{maximum}(l, y)$ when y is a proper upper bound of all elements in l , even if y is not itself an element of l). With the addition of the corule, we restrict the set of derivable judgments $\text{maximum}(l, x)$ to

those also admitting a finite-depth derivation using one of the two axioms, imposing that x must be an element of l . $\lrcorner$

Because of their interpretation, generalized inference systems are convenient to define mixed safety/liveness properties: safety properties are usually based on an invariance argument and can be naturally captured by the (coinductively interpreted) rules of the inference system; liveness properties are usually based on a well-foundedness argument and can be naturally captured by the (inductively interpreted) rules and corules. We will use generalized inference systems to provide compact definitions of fair subtyping (Section 3.2) and of the typing rules (Section 6). The reader interested in the metatheory of generalized inference systems may refer to [Ancona et al. \[2017\]](#) and [Dagnino \[2019\]](#).

3 SESSION TYPES

3.1 Syntax and Semantics

We use $l, a, b, \dots$ to denote the elements of a given set $\mathcal{L}$ of *labels* which may include values with a specific interpretation such as booleans, natural numbers, and so forth. A session type describes the communication protocol that takes place over a channel, namely the allowed sequences of input/output actions performed by a process on that channel. We use *polarities* $p \in \{?, !\}$ to distinguish *input actions* (?) from *output actions* (!) and we write $p^\perp$ for the *opposite* or *dual* polarity of p so that $?\perp = !$ and $!\perp = ?$. Session types are the possibly infinite, regular trees [[Courcelle 1983](#)] coinductively generated by the grammar below:

$$\text{Session type} \quad S, T ::= p \text{ end} \mid pS.T \mid p\{l_i : S_i\}_{i \in I}$$

Session types of the form $p \text{ end}$ describe channels used for exchanging a session termination signal and on which no further communication takes place. Session types of the form $pS.T$ describe channels used for exchanging another channel of type S and then according to T . Finally, session types of the form $p\{l_i : S_i\}_{i \in I}$ describe channels used for exchanging a label l_k and then according to S_k . Session types of the form $?\{l_i : S_i\}_{i \in I}$ and $!\{l_i : S_i\}_{i \in I}$ are sometimes referred to as *external* and *internal* choices respectively, to emphasize that the label being received or sent is always chosen by the sender process. In a session type $p\{l_i : S_i\}_{i \in I}$ we assume that I is not empty and that $i \neq j$ implies $l_i \neq l_j$ for every $i, j \in I$. Note that I is not necessarily finite, although regularity implies that there must be finitely many *distinct* S_i .

To improve readability we abbreviate $p\{l : S\}$ (when the choice is trivial) as $p!S$ and we define two partial operations $+$ and $\oplus$ such that

$$?\{l : S_l\}_{l \in A} + ?\{l : S_l\}_{l \in B} = ?\{l : S_l\}_{l \in A \cup B} \quad \text{and} \quad !\{l : S_l\}_{l \in A} \oplus !\{l : S_l\}_{l \in B} = !\{l : S_l\}_{l \in A \cup B} \quad (2)$$

when $A, B \neq \emptyset$ and $A \cap B = \emptyset$. We use U and V in addition to S and T to range over session types. Hereafter we specify possibly infinite session types by means of equations $S = \dots$ where the right hand side of the equation may contain guarded occurrences of the metavariable S . Guardedness guarantees that a session type S satisfying the equation exists and is unique [[Courcelle 1983](#)].

We equip session types with a *labeled transition system* (LTS) that allows us to describe, at the type level, the sequences of actions performed by a process on a channel. We distinguish two kinds of transitions: *unobservable transitions* $S \longrightarrow T$ are made autonomously by the process; *observable transitions* $S \xrightarrow{\alpha} T$ are made by the process in cooperation with the one it is interacting with through the channel. The label α describes the kind of interaction and has either the form pS (indicating the exchange of a channel of type S) or the form pl (indicating the exchange of label l). The polarity p indicates whether the message is received (?) or sent (!). The LTS is defined below:

$$?S.T \xrightarrow{?S} T \quad !S.T \xrightarrow{!S} T \quad ?\{l : S_l\}_{l \in A} \xrightarrow{?l} S_l \quad !l.S \oplus T \longrightarrow !l.S \quad !l.S \xrightarrow{!l} S$$

Note the different behaviors described by session types of the form $p\{l_i : S_i\}_{i \in I}$ depending on the polarity p . A process using a channel of type $?\{l_i : S_i\}_{i \in I}$ performs an observable transition for each of the labels l_i it is willing to receive. On the contrary, a process using a channel of type $!\{l_i : S_i\}_{i \in I}$ first *chooses* a particular label $l = l_k$ for some $k \in I$ (this choice is internal to the process and is therefore unobservable) and then *sends* the label l . As an example, the chain of transitions

$$!l.S \oplus T \longrightarrow !l.S \xrightarrow{!l} S$$

models a process that first chooses and then sends the label l . The choice of the label is irrevocable and not negotiable with the receiver process. Note that, according to the definition of $\oplus$ (cf. Eq. (2)), T must be an internal choice of labels different from l , hence $!l.S \oplus T$ is a non-trivial choice among two or more labels. Also, $!l.S$ admits no further unobservable transitions.

In the following we write $\Longrightarrow$ for the reflexive, transitive closure of $\longrightarrow$ and $\xRightarrow{\alpha}$ for the composition $\Longrightarrow \xrightarrow{\alpha}$. We extend transitions to strings of actions so that $\xRightarrow{\alpha_1 \cdots \alpha_n}$ stands for the composition $\xRightarrow{\alpha_1} \cdots \xRightarrow{\alpha_n}$. We let φ, ψ range over strings of actions, we write ε for the empty string of actions, $\leq$ for the usual prefix relation on strings, $S \xRightarrow{\varphi} T$ if $S \xRightarrow{\varphi} T$ for some T and $S \not\xRightarrow{\varphi}$ if not $S \xRightarrow{\varphi}$.

Definition 3.1 (paths of a session type). We say that φ is a *path* of S if $S \xRightarrow{\varphi}$. We write $\text{paths}(S)$ for the (prefix-closed) set of paths of S , that is $\text{paths}(S) \stackrel{\text{def}}{=} \{\varphi \mid S \xRightarrow{\varphi}\}$.

Note the difference between the relations $\Longrightarrow$ and $\xRightarrow{\varepsilon}$. The former relation entails zero or more unobservable transitions, whereas the latter relation entails no transitions at all. For example, we have $!l.S \oplus T \Longrightarrow !l.S$ but $!l.S \oplus T \not\xRightarrow{\varepsilon} !l.S$. As a consequence, the session type T on the right hand side of a relation $S \xRightarrow{\varphi} T$ is guaranteed to be a *subtree* of S . This property is useful to uniquely identify a particular subtree of S by means of a path φ .

Definition 3.2 (residual of a session type). The *residual* of a session type S with respect to a path $\varphi \in \text{paths}(S)$, denoted by $S(\varphi)$, is the unique session type T such that $S \xRightarrow{\varphi} T$.

Remark 1. Recall that regular trees are made of finitely many *distinct* subtrees [Courcelle 1983]. In particular, the set of all the residuals of S , namely $\{S(\varphi) \mid \exists \varphi \in \text{paths}(S)\}$, is always finite for every S . However, this set does not necessarily include *all* of the subtrees of S , since session types U occurring in prefixes of the form $pU.T$ are not reachable through paths. $\dashv$

The family of session types describing protocols that can always eventually terminate are particularly important in this work. We say that a session type with this property is *bounded*.

Definition 3.3 (bounded session type). We say that S is *bounded* if, for every $\varphi \in \text{paths}(S)$, there exist ψ and p such that $S(\varphi\psi) = p \text{ end}$.

That is, every path of a bounded session type can be extended to a maximal one after which the protocol is ended. For example, the session type $S = !\text{add}.S \oplus !\text{pay}.\text{end}$ is bounded whereas $T = !\text{add}.T$ is not. Note the difference between *finite* and bounded session types: every finite session type is bounded, but not every bounded session type is finite as illustrated by T above.

3.2 Fair Subtyping

The key ingredient of our type system is *fair subtyping*, denoted by the symbol $\leq$. The basic intuition underlying the subtyping relation $S \leq T$ is given by the Liskov substitution principle [Liskov and Wing 1994]: if S is a subtype of T , channels of type S can be “safely” used where channels of type T

Table 1. Fair subtyping.

$\frac{[F\text{-CONVERGE}]}{\forall \varphi \in \text{paths}(S) \setminus \text{paths}(T) : \exists \psi \leq \varphi, l \in \mathcal{L} : S(\psi!l) \leq T(\psi!l)}$ $S \leq T$			$\frac{[F\text{-END}]}{p \text{ end} \leq p \text{ end}}$
$\frac{[F\text{-CHANNEL}]}{S \leq T}$ $\frac{pU.S \leq pU.T}{S \leq T}$	$\frac{[F\text{-LABEL-IN}]}{S_i \leq T_i \ (i \in I)}$ $\frac{I \subseteq J}{\{l_i : S_i\}_{i \in I} \leq \{l_j : T_j\}_{j \in J}}$	$\frac{[F\text{-LABEL-OUT}]}{S_j \leq T_j \ (j \in J)}$ $\frac{J \subseteq I}{\{l_i : S_i\}_{i \in I} \leq \{l_j : T_j\}_{j \in J}}$	

are expected. We quote “safely” to stress that, in our work, subtyping is meant not only to preserve safety but also fair termination, which is a liveness property.

Fair subtyping is defined by the generalized interpretation of the inference system in Table 1. In addition to $\leq$, we write $\leq_{\text{ind}}$ for the relation defined by the *inductive* interpretation of the same inference system and $\leq_{\text{coind}}$ for the relation defined by the *coinductive* interpretation of the inference system in Table 1 excluding the corule [F-CONVERGE]. See Section 2 for the meaning of generalized and (co)inductive interpretations of a given set of inference rules.

To get a sense of fair subtyping, think of the relation $!a.S \oplus !b.T \leq !a.S$ saying that a channel of type $!a.S \oplus !b.T$ can be safely used where a channel of type $!a.S$ is expected. Indeed, a (well-typed) process that owns a channel of type $!a.S$ will use it for sending an a label and then according to S . This behavior is also allowed by the protocol $!a.S \oplus !b.T$, hence safety is preserved by the substitution. Dually, the relation $?a.S \leq ?a.S + ?b.T$ holds because a (well-typed) process that uses a channel of type $?a.S + ?b.T$ may receive either an a label or a b label, hence safety is preserved if the channel is replaced by another one of type $?a.S$ from which only an a label can be received.

The relation $\leq_{\text{coind}}$ – which we dub *unfair subtyping* – is essentially the same relation defined by Gay and Hole [2005], except that $\leq_{\text{coind}}$ is *invariant* for higher-order session types (cf. [F-CHANNEL]). The reason for this restriction is that we need to account for *all* usages of fair subtyping and make sure that they are done in suitably identified regions of a process. Allowing variant subtyping for higher-order session types may introduce “undetected” usages of fair subtyping that can compromise fair termination. We will dedicate Section 7 to analyzing this problem more in detail. The subtle difference between fair and unfair subtyping is due to the corule [F-CONVERGE]. Since this corule is somewhat obscure, we explain it gradually starting with the following observations:

- (1) Recall from Section 2 that $S \leq T$ implies $S \leq_{\text{coind}} T$ and $S \leq_{\text{ind}} T$. Hence, $\leq$ is a refinement of $\leq_{\text{coind}}$ such that, for each pair of related session types S and T , there exists a *finite-depth* derivation tree for the judgment $S \leq T$ using the rules and possibly the corule [F-CONVERGE].
- (2) When $S \leq_{\text{coind}} T$ holds, it is not possible to establish a general correlation between $\text{paths}(S)$ and $\text{paths}(T)$. Indeed, [F-LABEL-IN] entails that some paths of T are not present in S if $I \subsetneq J$ and [F-LABEL-OUT] entails that some paths of S are not present in T if $J \subsetneq I$.
- (3) The judgment $S \leq T$ is trivially derivable using [F-CONVERGE] if the path inclusion relation $\text{paths}(S) \subseteq \text{paths}(T)$ holds. Since [F-LABEL-OUT] is the only rule that allows T to have fewer paths than S , we deduce that [F-CONVERGE] limits (but does not always forbid) applications of [F-LABEL-OUT] when $J \subsetneq I$.
- (4) In general [F-CONVERGE] requires that, whenever a path φ of S is no longer present in T , it must be possible to find a prefix ψ of φ and an output $!l$ shared by both S and T such that $S(\psi!l)$ and $T(\psi!l)$ are one step closer to the region of S and T where path inclusion holds.

The reason why path inclusion plays such an important role in the definition of fair subtyping is that a process using a channel x of type T keeps using x according to T even if x is replaced by another channel of type $S \leq T$, without even realizing that the replacement has taken place. After all, this is what the “safe substitution principle” is based on. As a consequence, none of the paths in S that have disappeared in T will be offered to the process at the other end of the session x . If there are “too few” paths in T compared to S , then the replacement might compromise the termination of the process at the other end of the session, should it crucially rely on those paths to terminate. When $S \leq T$ (and therefore $S \leq_{\text{ind}} T$) holds, the corule [F-CONVERGE] makes sure that the process using x of type S believing that x has type T is always at *finite distance* from the region where path inclusion between (some subtrees of) S and (the corresponding subtrees of) T holds. Moreover, this region is always reachable by means of *output actions* (those $!l$ mentioned in [F-CONVERGE]) which are performed actively by the process using x . In other words, the process using x is always able, in a finite amount of time and relying on choices and actions it can perform autonomously, to steer the interaction towards a region of the protocol where path inclusion holds, hence where a common path to session termination is guaranteed to exist.

As we will see in the examples below, there might be infinitely many paths in $\text{paths}(S) \setminus \text{paths}(T)$ even if $S \leq T$. Nonetheless, [F-CONVERGE] is guaranteed to have finitely many premises because of the regularity of session types (cf. Remark 1).

Example 3.4 (acquirer protocols). Consider the session types $S = !\text{add}.S \oplus !\text{pay}.\text{end}$ and $T = !\text{add}.(!\text{add}.T \oplus !\text{pay}.\text{end})$ which might describe the protocols of potential acquirers in Eq. (1). The session type S describes an acquirer that purchases an arbitrary number of items, whereas the session type T describes an acquirer that always purchases an odd number of items.

Let us prove that $S \leq T$ holds. To this aim, we have to find a possibly infinite derivation tree for $S \leq T$ using the rules of Table 1 and, for each judgment $S' \leq T'$ in this derivation, also a finite derivation tree using the rules and the corule [F-CONVERGE]. The (infinite) derivation tree

$$\frac{\begin{array}{c} \vdots \\ S \leq T \end{array} \quad \frac{}{!\text{end} \leq !\text{end}} \text{[F-END]}}{\frac{S \leq !\text{add}.T \oplus !\text{pay}.\text{end}}{S \leq T} \text{[F-LABEL-OUT]}} \text{[F-LABEL-OUT]}$$

proves that $S \leq_{\text{coind}} T$ holds. There are three judgments occurring in this tree, namely $S \leq T$, $S \leq !\text{add}.T \oplus !\text{pay}.\text{end}$ and $!\text{end} \leq !\text{end}$, for which the following (finite) derivation trees can be obtained using [F-CONVERGE] and [F-END]:

$$\frac{\frac{}{!\text{end} \leq !\text{end}} \text{[F-END]}}{S \leq T} \text{[F-CONVERGE]} \quad \frac{\frac{}{!\text{end} \leq !\text{end}} \text{[F-END]}}{S \leq !\text{add}.T \oplus !\text{pay}.\text{end}} \text{[F-CONVERGE]}$$

More generally, if $T_{n \in \mathbb{N}}$ is the family of session types such that

$$T_n = \underbrace{!\text{add} \dots !\text{add}}_n . (!\text{add}.T_n \oplus !\text{pay}.\text{end})$$

where each T_n has an initial sequence of n $!\text{add}$ prefixes, it is possible to obtain similar derivation trees for $S \leq T_n$ for all $n \in \mathbb{N}$. However, if we consider the session type $T_\infty = !\text{add}.T_\infty$, which is somehow the limit of the succession $T_{n \in \mathbb{N}}$, we see that $S \leq_{\text{coind}} T_\infty$ holds but $S \leq T_\infty$ does not. Indeed, each $\varphi \in \text{paths}(S) \setminus \text{paths}(T_\infty)$ has the form $(!\text{add})^k !\text{pay}$ for some $k \in \mathbb{N}$ and there is no prefix ψ of φ and action $!l$ such that $S(\psi!l) \leq_{\text{ind}} T_\infty(\psi!l)$. The difference between T_n and T_∞ is that an acquirer behaving as T_n periodically has an opportunity of sending pay , which is essential for

the termination of the business, whereas an acquirer behaving as T_∞ keeps sending **add** forever. With this behavior, safety is preserved but fair termination is not. $\lrcorner$

Example 3.5 (random bit generator). In this example we see that there is no trivial correlation between fair subtyping and session type boundedness, contrarily to what the previous example might suggest. To this aim, imagine a service that generates random bits on demand. Its protocol could be described by the session type $S = ?\text{more}.(!0.S \oplus !1.S) + ?\text{stop}!\text{end}$ according to which the service sends a random bit if the client sends **more** and terminates if the client sends **stop**. Consider now a *fully biased* random bit generator that deterministically sends 0 on request. Its protocol is described by the session type $T = ?\text{more}!.0.T + ?\text{stop}!\text{end}$ and now we have that $S \leq_{\text{coind}} T$ holds whereas $S \leq T$ does not. The fact that fair subtyping does not hold can be shown with an argument similar to that used in Example 3.4, although we will provide a much easier proof in Section 3.3. The point to notice here is that T is bounded just like S is. Interestingly, we have $S \leq_{\text{coind}} T'$ and $S \not\leq T'$ also if T' describes a *partially biased* random bit generator, which deterministically sends two (or more) 0s in succession:

$$T' = ?\text{more}.(!0.(?\text{more}!.0.T' + ?\text{stop}!\text{end}) \oplus !1.T') + ?\text{stop}!\text{end}$$

In summary, the contravariance of label outputs allowed by [F-LABEL-OUT] may be constrained in non-trivial ways by [F-CONVERGE], depending on how input and output actions alternate. $\lrcorner$

We conclude this overview of fair subtyping by stating two notable properties of $\leq$.

PROPOSITION 3.6. (1) $\leq$ is a preorder and (2) if S, T are finite, then $S \leq T$ if and only if $S \leq_{\text{coind}} T$.

Concerning Item 1, while the reflexivity of $\leq$ is trivial and the transitivity or $\leq_{\text{coind}}$ is folklore [Gay and Hole 2005], the proof of transitivity of $\leq$ is made complex by the corule [F-CONVERGE]. In fact, the only proof we know of this fact relies on a semantic characterization of $\leq$ like the one we describe in Section 3.3. Item 2 shows that fair and unfair subtyping coincide for finite session types, hence the two relations only differ when infinite session types are considered. In particular, it can be shown that the corule [F-CONVERGE] is admissible if only finite session types are considered.

3.3 Compatibility

In this section we develop the notion of *session type compatibility*, which is instrumental to our theory for several reasons. First of all, compatibility is the relation assuring that the two endpoints of a session are used in complementary ways and that the amount of work that is necessary to terminate the session is finite (Section 3.4). This amount contributes to the definition of a *measure* at the level of processes and is a key element in the soundness proof of the type system (Section 6). Also, we use compatibility to formalize the relationship between fair subtyping (Section 3.2), the standard notion of *duality* for session types (Section 3.5) and the notion of *fair termination* found in the literature (Section 3.6). Finally, compatibility provides the semantic grounds to justify the corule [F-CONVERGE] in the definition of fair subtyping and consequently the technical machinery that we use to prove that fair subtyping is a preorder (Item 1 of Proposition 3.6).

Intuitively, S and T are *compatible* when they entail a “correct interaction” between the two processes that use the peer endpoints of a session, one of type T and the other of type S . If we use a term of the form $S \mid T$ to describe the session as a whole, with the two interacting processes behaving as S and T , then we can formalize their interaction (at the type level) using the LTS for session types and the reduction rules

$$\frac{S \longrightarrow S'}{S \mid T \longrightarrow S' \mid T} \quad \frac{T \longrightarrow T'}{S \mid T \longrightarrow S \mid T'} \quad \frac{S \xrightarrow{\alpha^+} S' \quad T \xrightarrow{\alpha} T'}{S \mid T \longrightarrow S' \mid T'}$$

where we write $\alpha^\perp$ for the dual action of α , obtained by changing the polarity of α with the opposite one. A reduction occurs whenever one of the connected processes performs an unobservable transition or when the two processes exchange a message by proposing complementary actions. As usual, we let $\Longrightarrow$ stand for the reflexive, transitive closure of $\longrightarrow$ and we write $S \mid T \longrightarrow$ if there are no S' and T' such that $S \mid T \longrightarrow S' \mid T'$.

Session type compatibility is the property saying that every finite interaction between two peer processes over a session can always be extended so as to successfully terminate the session.

Definition 3.7 (session type compatibility). We say that S and T are *compatible*, notation $S \sim T$, if $S \mid T \Longrightarrow S' \mid T' \implies S' \mid T' \Longrightarrow p^\perp \text{ end} \mid p \text{ end}$ for some p .

Note that compatibility implies the absence of communication errors, whereby a channel of unexpected type or an unexpected label is exchanged. Indeed, both $!a.S \mid ?b.T$ and $!U.S \mid ?V.T$ are stuck if $a \neq b$ and $U \neq V$. Compatibility also implies progress, for a session where both peers simultaneously attempt at receiving or sending a message is (or becomes) stuck.

Example 3.8. Consider the session types defined in Example 3.4 as well as $U = ?\text{add}.U + ?\text{pay}.\text{end}$. It is easy to see that U is compatible with both S and T_1 as well as with all the T_n for $n \in \mathbb{N}$. However, U is incompatible with T_∞ – despite the fact that U and T_∞ may interact forever without getting stuck – because $U \mid T_\infty$ cannot reach the state $?\text{end} \mid !\text{end}$. In other words, U provides a semantic justification for $S \not\leq T_\infty$: a business that behaves according to U can (fairly) terminate if it interacts with any acquirer that behaves according to S or any of the T_n , but not if it interacts with an acquirer that behaves according to T_∞ . $\dashv$

The next two results formalize the tight relationship between compatibility and fair subtyping. First of all, fair subtyping preserves compatibility.

THEOREM 3.9. *If $S \leq T$, then $U \sim S$ implies $U \sim T$ for every U .*

Theorem 3.9 seems to reverse the direction of the substitution principle that we mentioned when introducing fair subtyping (Section 3.2). The contradiction is only apparent, however, and is resolved by observing that Liskov’s principle speaks of right-to-left substitutability of *values/channels* whereas Theorem 3.9 speaks of left-to-right substitutability of *behaviors/processes*. Gay [2016] discusses more in detail these two different yet related viewpoints.

Fair subtyping is also the *coarsest* subtyping relation between (bounded) session types that preserves compatibility. More precisely:

THEOREM 3.10. *If S is bounded and $U \sim S$ implies $U \sim T$ for every U , then $S \leq T$.*

Theorem 3.10 shows that the corule [F-CONVERGE] is a *necessary condition* to turn unfair subtyping into a compatibility-preserving subtyping relation, at least when bounded session types are related by fair subtyping. In other words, if $S \leq_{\text{coind}} T$ and $S \not\leq T$, then it is possible to find U that is compatible with S but not with T , as we have done in Examples 3.8 and 3.11.

We can combine Theorem 3.9 and Theorem 3.10 to prove the transitivity of $\leq$ (among bounded session types). Indeed, suppose that $S \leq U$ and $U \leq T$ hold, where S is bounded, and consider an arbitrary V that is compatible with S . By Theorem 3.9 we deduce that V is compatible with U and therefore with T . By Theorem 3.10 we conclude $S \leq T$.

Example 3.11. Consider once again the session types S and T defined in Example 3.5, which are not related by fair subtyping, and let $U = !\text{more}.(?0.U + ?1.\text{stop}.\text{end})$. Note that U stops the interaction as soon as it receives a 1 from the random bit generator. For this reason, we have $U \sim S$ but $U \not\sim T$ since T never sends 1 . By Theorem 3.9, we deduce $S \not\leq T$ as we had already argued in Example 3.5. We can use a similar reasoning to show that $S \not\leq T'$, except that the witness behavior

that distinguishes S from T' is slightly more involved. The idea is to design a session type V that ends as soon as it receives a 1 immediately after it has received a 0 :

$$V = !\text{more}.(?0.\text{more}.(?0.V + ?1.\text{stop}.\text{end}) + ?1.V)$$

The proof of Theorem 3.10 is based on an effective construction of a discriminating session type (such as U or V above) that is compatible with S but not with T whenever $S \leq_{\text{coind}} T$ holds but $S \leq_{\text{ind}} T$ does not. $\dashv$

Remark 2. The reason why Theorem 3.10 does *not* hold in general, but only when S is bounded, is that the notion of session type compatibility that we consider (Definition 3.7) induces a large family of session types that are semantically equivalent (in the sense that they are not compatible with any other session type) but syntactically unrelated. As an example, the session types $S = ?\text{add}.S$ and $T = !\text{pay}.T$ are incompatible with any other session type (including themselves) simply because they do not contain an end leaf. In this case it is trivially true that any session type compatible with S is also compatible with T , but S and T cannot be related using the definition of $\leq$ as it stands. To make the correspondence between fair subtyping and compatibility preservation exact it is necessary to adopt a slightly different notion of session type compatibility that is biased towards the successful termination of one of the two session participants [Ciccone and Padovani 2021b; Padovani 2013]. This one-sided compatibility does not capture exactly the notion of “correct session termination” as we intend it, according to which *both* participants are required to successfully terminate, but is the one used in the proof that $\leq$ is transitive in general. $\dashv$

3.4 Rank of a Session

In the soundness proof of our type system we need to quantify the amount of work required to terminate a particular session in which one process behaves as S and the other as T . To this aim, we define the rank of this session as the smallest number of interactions that lead $S \mid T$ to termination.

Definition 3.12 (rank). The rank of S and T , written $\|S, T\|$, is the element of $\mathbb{N} \cup \{\infty\}$ defined as

$$\|S, T\| \stackrel{\text{def}}{=} \min\{1 + |\varphi| \mid \exists \varphi, p : S \xrightarrow{\varphi^\perp} p^\perp \text{end}, T \xrightarrow{\varphi} p \text{end}\}$$

where $|\varphi|$ denotes the length of φ , $\varphi^\perp$ is the string of actions obtained by dualizing all the actions in φ , and we postulate that $\min \emptyset = \infty$.

As we will see in Section 4, our process calculus requires an explicit message exchange for closing a session. This is the reason why we add 1 to the length of all paths that lead S and T to termination, so that the rank of S and T measures the actual number of synchronizations that are necessary to terminate the session. Note that the rank $\|S, T\|$ is generally unrelated to the lengths of the shortest paths of S and T that lead to termination. For example, if we take $S = ?a.!c.?a.?end + ?b.?end$ and $T = !a.(?c.!a.!end + ?d.!end)$ we see that the shortest path φ such that $S(\varphi) = ?end$ is $?b$ of length 1 and the shortest path ψ such that $T(\psi) = !end$ is $!a?d$ of length 2, but $\|S, T\| = 4$.

The rank of two compatible session types is always finite and varies in agreement with subtyping:

THEOREM 3.13. *If $U \sim S$, then (1) $\|U, S\| \in \mathbb{N}$ and (2) $S \leq_{\text{coind}} T$ implies $\|U, S\| \leq \|U, T\|$.*

Theorem 3.13 shows that every usage of (fair) subtyping may increase the amount of work that is necessary to terminate a session (Item 2), although such amount is guaranteed to remain finite as long as compatibility is preserved (Item 1). This property justifies the adoption of fair subtyping over unfair subtyping, since fair subtyping preserves compatibility (Theorem 3.9) whereas unfair subtyping in general does not (Examples 3.8 and 3.11). Theorem 3.13 also suggests that the finiteness of the rank can be guaranteed only when fair subtyping is used *finitely many times*. For this reason,

we will have to be careful on *where* fair subtyping is used in the typing derivation of recursive processes to avoid that “too many” applications of fair subtyping end up having the same effect of unfair subtyping (Section 5.3).

3.5 Duality

In binary session type theories, the two endpoints of a session are associated with session types such that one is the *dual* of the other. The dual of a session type S has the same overall structure of S , but opposite polarities for the corresponding actions. Formally:

Definition 3.14. The *dual* of a session type S , written $S^\perp$, is corecursively defined by the equations

$$(p \text{ end})^\perp = p^\perp \text{ end} \quad (pS.T)^\perp = p^\perp S.T^\perp \quad (p\{l_i : S_i\}_{i \in I})^\perp = p^\perp \{l_i : S_i^\perp\}_{i \in I}$$

Although duality guarantees communication safety and progress, it does not imply compatibility in general. To see this, consider the session type T_∞ from Example 3.4 and note that

$$T_\infty^\perp \mid T_\infty \longrightarrow T_\infty^\perp \mid T_\infty \longrightarrow \dots \not\Rightarrow p^\perp \text{ end} \mid p \text{ end}$$

so two processes adhering to $T_\infty^\perp$ and T_∞ would be able to interact forever, but without hope of terminating the session. This kind of interaction must be forbidden in our setting if we are interested in fairly terminating sessions, hence $T_\infty^\perp \not\sim T_\infty$. Still, there is a connection between duality, boundedness and compatibility that is fundamental in the proof of Theorem 3.10, as it guarantees that every bounded session type is compatible with at least another one, its dual.

THEOREM 3.15. $U^\perp \sim U$ if and only if U is bounded.

3.6 Fair Termination

In this section we relate compatibility with the notions of strong fairness and fair termination found in the literature [Apt et al. 1987; Francez 1986; van Glabbeek and Höfner 2019]. In general, fairness assumptions are made to rule out those infinite runs of a system that are considered unrealistic. In order to formulate strong fairness, we must therefore define a notion of “run” in our setting.

Definition 3.16 (run). A *run* of $S \mid T$ is a sequence of reductions

$$S \mid T \longrightarrow S_1 \mid T_1 \longrightarrow S_2 \mid T_2 \longrightarrow \dots$$

and it is *maximal* if either it is infinite or if it ends with a term $S_n \mid T_n$ such that $S_n \mid T_n \not\rightarrow$.

Among all possible runs, we identify the “fair” ones as those in which *reductions that are enabled infinitely often occur infinitely often*. In the taxonomy of fairness notions [Kwiatkowska 1989; van Glabbeek and Höfner 2019], this particular one is called *strong fairness*. Formally:

Definition 3.17 (fair run). A run π is *fair* if, for every $S \mid T$ that occurs infinitely often in π and every $S' \mid T'$ such that $S \mid T \longrightarrow S' \mid T'$, the reduction $S \mid T \longrightarrow S' \mid T'$ occurs infinitely often in π .

This notion of fairness is known to enjoy two properties that we use in our development: (1) every finite run is a fair run; (2) every finite run can be extended to a maximal fair run. The second property is considered to be an essential requirement for every fairness notion and is referred to in the literature as *machine closure* [Lamport 2000] or *feasibility* [Apt et al. 1987; van Glabbeek and Höfner 2019]. We can now define what it means for a session to be fairly terminating:

Definition 3.18 (fair termination). We say that a session described by the pair of session types $S \mid T$ *fairly terminates* if all of its maximal fair runs are finite.

Example 3.19. Consider T and T_∞ from Example 3.4 and the session type $U = ?\text{add}.U + ?\text{pay}.\text{end}$. We can depict all the runs of $U \mid T$ as the infinite tree

$$\begin{array}{c} U \mid T \longrightarrow U \mid !\text{add}.T \oplus !\text{pay}.\text{end} \longrightarrow U \mid !\text{add}.T \longrightarrow U \mid T \longrightarrow \dots \\ \downarrow \\ U \mid !\text{pay}.\text{end} \longrightarrow ?\text{end} \mid !\text{end} \end{array}$$

where we observe that every run of $U \mid T$ ending in $?\text{end} \mid !\text{end}$ is maximal, finite and therefore fair, whereas the only infinite run of $U \mid T$ is unfair since the reduction $U \mid !\text{add}.T \oplus !\text{pay}.\text{end} \longrightarrow U \mid !\text{pay}.\text{end}$ is infinitely often enabled but never performed. On the other hand, $U \mid T_\infty$ has only one maximal run $U \mid T_\infty \longrightarrow U \mid T_\infty \longrightarrow \dots$, which is infinite and fair. In summary, the session $U \mid T$ is fairly terminating, whereas the session $U \mid T_\infty$ is not. $\dashv$

We can now establish the tight relationship between the compatibility of S and T and fair termination of $S \mid T$ without residual pending communications.

THEOREM 3.20. *For every S and T we have that $S \sim T$ if and only if every maximal fair run of $S \mid T$ is finite and ends with $p^\perp \text{end} \mid p \text{end}$ for some p .*

The “only if” part of Theorem 3.20 (without the requirement that the final term of the fair run has the form $p^\perp \text{end} \mid p \text{end}$) is known as *liveness enhancing property* of the fairness assumption [Apt et al. 1987; van Glabbeek and Höfner 2019]. It shows that the fairness assumption affects the liveness properties that can be proved: some liveness properties (e.g. termination) do not hold in general (there exist infinite runs) but they do hold if the unfair runs are ruled out (all fair runs are finite).

4 LANGUAGE SYNTAX AND SEMANTICS

The syntax of processes makes use of an infinite set of *channel names*, ranged over by x, y and z , and of a finite set of *process names*, ranged over by A, B and C . Hereafter, we use $\bar{x}$ to denote a possibly empty tuple of names, extending the same notation to other entities. A *program* $\mathcal{P}$ is a finite set $\{A_i(\bar{x}_i) \triangleq P_i\}_{i \in I}$ of *definitions* where each P_i is a *process* generated by the grammar below:

Process	$P, Q ::=$	done	termination		$A(\bar{x})$	invocation
		wait $x.P$	signal input		close x	signal output
		$x?(y).P$	channel input		$x!y.P$	channel output
		$xp\{l_i : P_i\}_{i \in I}$	label input/output		$P \oplus_k Q$	choice
		$(x)(P \mid Q)$	session		$[x]P$	cast

The process **done** is terminated and performs no action. The invocation $A(\bar{x})$ behaves as P if $A(\bar{x}) \triangleq P$ is the definition of A . When $\bar{x}$ is empty, we write A and $A \triangleq P$ instead of $A(\langle \rangle)$ and $A(\langle \rangle) \triangleq P$. The process **wait** $x.P$ waits for a signal from channel x indicating that the session x is being closed and then continues as P . The process **close** x sends the termination signal on x . The process $x?(y).P$ receives a channel y from channel x and then continues as P . Dually, $x!y.P$ sends y on x and then continues as P . The process $xp\{l_i : P_i\}_{i \in I}$ exchanges a label l_i on channel x and then continues as P_i . As for session types, we assume that the set I in these forms is always non-empty and that $i \neq j$ implies $l_i \neq l_j$ for every $i, j \in I$. Also, we write $xpl_i.P_i$ instead of $xp\{l_i : P_i\}_{i \in I}$ when I is a singleton $\{i\}$. A non-deterministic choice $P_1 \oplus_k P_2$ reduces to either P_1 or P_2 . The annotation $k \in \{1, 2\}$ has no operational meaning, it is only used to record that P_k leads to the termination of the process and is omitted when irrelevant. A session $(x)(P \mid Q)$ is the parallel composition of P and Q connected by x . Finally, a *cast* $[x]P$ behaves exactly as P . This form simply records the fact that the type of x is subject to an application of fair subtyping in the typing derivation for P . As we will see in Sections 5.3 and 6, we use this form to precisely account for all places in (the typing derivation of) a process where fair subtyping is used. Occasionally we write $[x_1 \dots x_n]P$ for $[x_1] \dots [x_n]P$.

Table 2. Structural pre-congruence and reduction of processes.

[S-PAR-COMM]	$(x)(P \mid Q) \leqslant (x)(Q \mid P)$	
[S-PAR-ASSOC]	$(x)(P \mid (y)(Q \mid R)) \leqslant (y)((x)(P \mid Q) \mid R)$	if $x \in \text{fn}(Q)$
[S-CAST-COMM]	$\lceil x \rceil \lceil y \rceil P \leqslant \lceil y \rceil \lceil x \rceil P$	
[S-CAST-NEW]	$(x)(\lceil x \rceil P \mid Q) \leqslant (x)(P \mid Q)$	
[S-CAST-SWAP]	$(x)(\lceil y \rceil P \mid Q) \leqslant \lceil y \rceil (x)(P \mid Q)$	if $x \neq y$
[S-CALL]	$A(\bar{x}) \leqslant P$	if $A(\bar{x}) \triangleq P$
[R-CHOICE]	$P_1 \oplus P_2 \longrightarrow P_k$	if $k \in \{1, 2\}$
[R-SIGNAL]	$(x)(\text{close } x \mid \text{wait } x.P) \longrightarrow P$	
[R-CHANNEL]	$(x)(x!y.P \mid x?(y).Q) \longrightarrow (x)(P \mid Q)$	
[R-PICK]	$(x)(x!\{l_i : P_i\}_{i \in I} \mid Q) \longrightarrow (x)(x!l_k.P_k \mid Q)$	if $k \in I$ and $ I > 1$
[R-LABEL]	$(x)(x!l_k.P \mid x?\{l_i : Q_i\}_{i \in I}) \longrightarrow (x)(P \mid Q_k)$	if $k \in I$
[R-PAR]	$(x)(P \mid R) \longrightarrow (x)(Q \mid R)$	if $P \longrightarrow Q$
[R-CAST]	$\lceil x \rceil P \longrightarrow \lceil x \rceil Q$	if $P \longrightarrow Q$
[R-STRUCT]	$P \longrightarrow Q$	if $P \leqslant P' \longrightarrow Q' \leqslant Q$

The only binders are channel inputs $x?(y).P$ and sessions $(x)(P \mid Q)$. We write $\text{fn}(P)$ and $\text{bn}(P)$ for the sets of free and bound channel names occurring in P and we identify processes modulo renaming of bound names. The program $\mathcal{P}$ that provides the meaning to the process names occurring in processes is often left implicit. Sometimes we write a process definition $A(\bar{x}) \triangleq P$ as a proposition or side condition, intending that such definition is part of the implicit program $\mathcal{P}$.

Example 4.1. Let us revisit and complete the example we sketched in Section 1. We can model the whole system as the following set of process definitions:

$\text{Main} \triangleq (y)((x)(\lceil x \rceil A(x) \mid B(x, y)) \mid C(y)) \quad B(x, y) \triangleq x?\{\text{add} : B(x, y), \text{pay} : \text{wait } x.y!\text{ship.close } y\}$
 $A(x) \triangleq x!\text{add}.x!\{\text{add} : A(x), \text{pay} : \text{close } x\} \quad C(y) \triangleq y?\text{ship.wait } y.\text{done}$

Note that the acquirer deterministically sends **add** to the business as the first message, whereas it chooses among **add** and **pay** every other interaction. After the acquirer has sent **pay**, it closes the session x with the business B . At this point, the business sends **ship** to the carrier C and closes the session y . The cast $\lceil x \rceil$ before the invocation of $A(x)$ in Main is meant to account for the mismatch between the behavior of the acquirer, which always adds an odd number of items to the cart, and that of the business, which accepts any number of items added to the shopping cart. $\dashv$

The operational semantics of processes is defined using a structural pre-congruence relation $\leqslant$ and a reduction relation $\longrightarrow$, both of which are defined in Table 2 and described hereafter. Rules [S-PAR-COMM] and [S-PAR-ASSOC] express the usual commutativity and associativity of parallel composition. In the case of [S-PAR-ASSOC], the side condition $x \in \text{fn}(Q)$ makes sure that the session $(x)(P \mid Q)$ we obtain on the right hand side does indeed connect P and Q through x . Also note that [S-PAR-ASSOC] only describes a right-to-left associativity of parallel composition and that left-to-right associativity is derivable. The remaining axioms are those that justify the use of a pre-congruence over a symmetric congruence relation. Since each usage of fair subtyping may increase the amount of work that is necessary to terminate a session (cf. Theorem 3.13), axiom [S-CAST-NEW] annihilates a cast on x nearby the binder for x , making sure that casts can only be removed and never added. Axioms [S-CAST-COMM] and [S-CAST-SWAP] are used to move casts closer to their binder so that they can be annihilated with [S-CAST-NEW]. Rule [S-CALL] unfolds process invocations to their definition.

The reduction rules are mostly unremarkable: **[R-CHOICE]** models the non-deterministic choice between alternative behaviors; **[R-PICK]** models a non-trivial choice among a set of labels to send; **[R-SIGNAL]**, **[R-LABEL]** and **[R-CHANNEL]** model synchronizations between a sender (on the left hand side of the parallel composition) and a receiver (on the right hand side of the parallel composition) with **[R-SIGNAL]** removing the binder of a closed session; **[R-PAR]**, **[R-CAST]** and **[R-STRUCT]** close reductions under parallel compositions, under casts and by structural pre-congruence. In the following we write $\Rightarrow$ for the reflexive, transitive closure of $\rightarrow$ and $\Rightarrow^+$ for $\Rightarrow \rightarrow$.

We can now define the property enforced by our type system.

Definition 4.2. We say that P is *fairly terminating* if $P \Rightarrow Q$ implies $Q \leq \text{done}$ or $Q \Rightarrow^+ \text{done}$.

Remark 3. The definitions of session type compatibility (Definition 3.7) and of fair process termination (Definition 4.2) are inspired to that of *successful computation* in fair testing theories [Natarajan and Cleaveland 1995; Rensink and Vogler 2007]. Rensink and Vogler [2007] show that these notions have a *built-in fairness assumption* that coincides with strong fairness, at least in the case of *finite-state processes* (in fact, Theorem 3.20 is a particular instance of this result for session types). But while defining fair runs for session types is doable with little effort (cf. Definition 3.17), the definition of fair runs for the π -calculus is much more involved [Biding and Compagnoni 2009; Cacciagrano et al. 2006, 2009; Kobayashi 2002]. Besides, none of the available definitions is directly applicable to our language since they are all based on choiceless versions of the π -calculus with replication instead of recursion. For these reasons, we adopt the formulation of fair process termination in Definition 4.2 for its appeal and simplicity: the reachability of **done** implies that every pending action (resp. open session) in Q may eventually be performed (resp. terminated). $\perp$

Example 4.3. With the definitions given in Example 4.1, it is easy to see that there is an infinite reduction sequence starting from *Main* in which the acquirer keeps adding items to the cart:

$$\begin{aligned} (y)((x)([x]A\langle x \rangle \mid B\langle x, y \rangle) \mid C\langle y \rangle) &\Rightarrow (y)((x)(x!\{\text{add} : A\langle x \rangle, \text{pay} : \text{close } x\} \mid B\langle x, y \rangle) \mid C\langle y \rangle) \\ &\rightarrow (y)((x)(x!\text{add}.A\langle x \rangle \mid B\langle x, y \rangle) \mid C\langle y \rangle) \\ &\Rightarrow (y)((x)(A\langle x \rangle \mid B\langle x, y \rangle) \mid C\langle y \rangle) \rightarrow \dots \end{aligned}$$

Nonetheless, *Main* is fairly terminating. For example, we have:

$$\begin{aligned} (y)((x)(x!\{\text{add} : A\langle x \rangle, \text{pay} : \text{close } x\} \mid B\langle x, y \rangle) \mid C\langle y \rangle) &\rightarrow (y)((x)(x!\text{pay.close } x \mid B\langle x, y \rangle) \mid C\langle y \rangle) \\ &\Rightarrow (y)((x)(\text{close } x \mid \text{wait } x.y!\text{ship.close } y) \mid C\langle y \rangle) \\ &\rightarrow (y)(y!\text{ship.close } y \mid C\langle y \rangle) \Rightarrow (y)(\text{close } y \mid \text{wait } y.\text{done}) \rightarrow \text{done} \end{aligned}$$

Note that in general it might be necessary for the acquirer to add one more item to the cart before it can send the payment to the business and the carrier receives a **ship** message. $\perp$

5 THE TYPE SYSTEM BY EXAMPLES

In this section we motivate, through a series of examples, the key properties enforced by the type system that, taken together, guarantee fair termination. There are two families of problems that can compromise fair termination. First of all, the process (or part thereof) may be unable to reduce further but is not **done**. In our model, this can happen for many reasons, for example: a process attempts at sending a label on a session that the receiver is not willing to accept; a process attempts at sending a termination signal when the receiver expects a channel; the processes at the two ends of the same session are both waiting for a message from that session. These are all examples of *safety violations*, which are prevented by any ordinary session type system. In this section we focus instead on *liveness violations*. Roughly speaking, liveness is violated when a process (or part thereof) engages an infinite computation that cannot possibly terminate. In Section 3.2 we have introduced

a fair subtyping relation that is liveness preserving but, as we will see in a moment, the adoption of fair subtyping alone is not enough to rule out all potential liveness violations. The type system must also enforce three properties that we call *action boundedness*, *session boundedness* and *cast boundedness* guaranteeing that the overall effort required to terminate the process is finite. In the rest of the section we describe informally these properties and we show that violating even just one of them may compromise fair process termination. In doing so, we assume that the reader has some familiarity with the basic features of session type systems, those that prevent the aforementioned safety violations. If not, it might be worth revisiting this section after reading Section 6.

5.1 Action Boundedness

We say that a process is *action bounded* if there is a finite upper bound to the number of actions it has to perform in order to terminate. An action-unbounded process cannot terminate. Compare

$$A \triangleq A \oplus \text{done} \quad \text{and} \quad B \triangleq B \oplus B \quad (3)$$

and observe that A may always reduce to **done**, whereas B can only reduce forever into itself. So A is action bounded whereas B is not. We consider a parallel composition action bounded if so are *both* processes composed in parallel.

Action boundedness is a necessary condition for (fair) process termination, hence the type system must guarantee that well-typed processes are action bounded. As we will see in Section 6, this can be easily achieved by means of *typing corules*. Besides, action boundedness carries along two welcome side effects. The first one is that degenerate process definitions such as $A \triangleq A$ are not action bounded and therefore are flagged as ill typed by the type system. This guarantees that finitely many unfoldings of recursive process invocations always suffice to expose some observable process behavior. The second is that action boundedness allows us to detect recursive processes that claim to use a channel in a certain way when in fact they never do so. As an example, compare

$$A(x, y) \triangleq x!\text{a}.A\langle x, y \rangle \oplus x!\text{b}.\text{close } x \quad \text{and} \quad B(x, y) \triangleq x!\text{a}.B\langle x, y \rangle$$

where $A\langle x, y \rangle$ is action bounded and $B\langle x, y \rangle$ is not. An ordinary session type system with coinductively interpreted typing rules would accept $B\langle x, y \rangle$ regardless of y 's type on the grounds that y occurs once in the body of B , hence it is “used” linearly. This is unfortunate, since y is not used in any meaningful way other than being passed as an argument of B . In A , the same linearity check promptly detects that y is not used along the path to **close** x that proves the boundedness of $A\langle x, y \rangle$.

5.2 Session Boundedness

We say that a process is *session bounded* if there is a finite upper bound to the number of sessions it has to create in order to terminate. It is easy to construct non-terminating processes by chaining together an infinite number of finite (or fairly terminating) sessions. For example, compare

$$A \triangleq (x)(\text{close } x \mid \text{wait } x.A) \oplus \text{done} \quad \text{and} \quad B_1 \triangleq (x)(\text{close } x \mid \text{wait } x.B_1) \quad (4)$$

where A always has a possibility to terminate without creating new sessions (it is session bounded) while B_1 does not (it is session unbounded). It could be argued that B_1 is already ruled out because it is not action bounded. Indeed, while the left-hand side of the parallel composition in B_1 is finite, the right hand side is not (recall that we require *both* sides of a parallel composition to admit a finite path to either **done** or **close** x). Below is a slightly more complex variation of B_1 that is action bounded and session unbounded. The trick is to have a finite branch on one side of the parallel composition matched by an infinite one on the other side:

$$B_2 \triangleq (x)(x!\{a : \text{close } x, b : \text{wait } x.B_2\} \mid x?\{a : \text{wait } x.B_2, b : \text{close } x\}) \quad (5)$$

Eq. (4) shows that a session bounded process like A may still create an unbounded number of sessions. Below is another example of session bounded process that creates unboundedly many *nested* sessions, such that the first session being created is also the last one being completed:

$$(x)(C\langle x \rangle \mid \text{wait } x.\text{done}) \quad \text{where} \quad C(x) \triangleq (y)(C\langle y \rangle \mid \text{wait } y.\text{close } x \oplus \text{close } x) \quad (6)$$

While both A and C may create an arbitrary number of sessions, they do not *have to* do so in order to terminate. This is what sets them apart from B_1 and B_2 .

5.3 Cast Boundedness

We say that a process is *cast bounded* if there is a finite upper bound to the number of casts it has to perform in order to terminate. Performing a cast means applying [S-CAST-NEW], which corresponds to a usage of fair subtyping. The reason why cast boundedness is fundamental is that the liveness-preserving property of fair subtyping holds as long as fair subtyping is used finitely many times. Conversely, infinitely many usages of fair subtyping may have the overall effect of a single usage of unfair subtyping (cf. Example 3.4). By “infinitely many usages” we mean usages of fair subtyping that occur within a loop in a recursive process. To illustrate the problem, let us consider the (non-terminating) process

$$(x)(A\langle x \rangle \mid B\langle x \rangle) \quad \text{where} \quad \begin{aligned} A(x) &\triangleq [x]x!\text{add}.A\langle x \rangle \\ B(x) &\triangleq x?\{\text{add} : B\langle x \rangle, \text{pay} : \text{wait } x.\text{done}\} \end{aligned} \quad (7)$$

and the session type $S = !\text{add}.S \oplus !\text{pay}.\text{end}$. It can be argued that the channel x is used according to S in $A(x)$ and according to $S^\perp$ in $B(x)$. Indeed, the structure of $B(x)$ matches perfectly that of $S^\perp$ and $x!\text{add}.A(x)$ uses x according to $!\text{add}.S$, which is a fair supertype of S accounted for by the cast $[x]$ in A . With this cast it is as if $A(x)$ promises to make a choice between sending **add** and sending **pay** at each iteration, but systematically favors **add** over **pay**. The overall effect of these unfulfilled promises is that the actual behavior of $A(x)$ over x is better described by the session type $T_\infty = !\text{add}.T_\infty$, which is *not* a fair supertype of S as we have seen in Examples 3.4 and 3.8.

Although $A(x)$ could be rejected on the grounds that it is not action bounded, it is possible to find an action-bounded (but slightly more involved) variation of $A(x)$ and $B(x)$ in Eq. (7) in which the same phenomenon occurs. With the definitions

$$\begin{aligned} A(x) &\triangleq [x]x!\text{more}.x?\{\text{more} : A\langle x \rangle, \text{stop} : \text{wait } x.\text{done}\} \\ B(x) &\triangleq x?\{\text{more} : [x]x!\text{more}.B\langle x \rangle, \text{stop} : \text{wait } x.\text{done}\} \end{aligned} \quad (8)$$

both $A(x)$ and $B(x)$ have a chance to continue or to terminate the session by sending either **more** or **stop**, except that they systematically favor **more** over **stop**. Now, if we consider the session type $S = !\text{more}.(? \text{more}.S + ? \text{stop}.\text{end}) \oplus !\text{stop}.\text{end}$, it can be argued that $A(x)$ uses x according to $S_A = !\text{more}.(? \text{more}.S + ? \text{stop}.\text{end})$, which is a fair supertype of S , and that $B(x)$ uses x according to $S_B = ? \text{more}.\text{more}.S^\perp + ? \text{stop}.\text{end}$, which is a fair supertype of $S^\perp$. The two casts in Eq. (8) account for the differences between S and S_A in $A(x)$ and between $S^\perp$ and S_B in $B(x)$, but they occur within loops along paths that lead to process termination, hence A and B are not cast bounded.

It is worth discussing one last attempt to work around the problem, by moving the casts outward from within $A(x)$ and $B(x)$, as in

$$(x)([x]A\langle x \rangle \mid [x]B\langle x \rangle) \quad \text{where} \quad \begin{aligned} A(x) &\triangleq x!\text{more}.x?\{\text{more} : A\langle x \rangle, \text{stop} : \text{wait } x.\text{done}\} \\ B(x) &\triangleq x?\{\text{more} : x!\text{more}.B\langle x \rangle, \text{stop} : \text{wait } x.\text{done}\} \end{aligned} \quad (9)$$

Now $A(x)$ uses x according to $T_A = !\text{more}.(? \text{more}.T_A + ? \text{stop}.\text{end})$ and $B(x)$ uses x according to $T_B = ? \text{more}.\text{more}.T_B + ? \text{stop}.\text{end}$, but while $S \leq_{\text{coind}} T_A$ and $S^\perp \leq_{\text{coind}} T_B$ both hold neither $S \leq T_A$ nor $S^\perp \leq T_B$ does. In summary, the non-terminating process in Eq. (9) is action bounded, session bounded and cast bounded, but it is typeable only using *unfair* subtyping.

Table 3. Typing rules.

$\frac{}{\emptyset \vdash^n \text{done}} \quad [\text{T-DONE}]$	$\frac{\Gamma \vdash^n P}{\Gamma, x : ?\text{end} \vdash^n \text{wait } x.P} \quad [\text{T-WAIT}]$	$\frac{}{x : !\text{end} \vdash^n \text{close } x} \quad [\text{T-CLOSE}]$	$\frac{\Gamma, x : S, y : T \vdash^n P}{\Gamma, x : ?T.S \vdash^n x?(y).P} \quad [\text{T-CHANNEL-IN}]$
$\frac{\Gamma, x : S \vdash^n P}{\Gamma, x : !T.S, y : T \vdash^n x!y.P} \quad [\text{T-CHANNEL-OUT}]$	$\frac{\Gamma, x : S_i \vdash^n P_i \ (i \in I)}{\Gamma, x : p\{l_i : S_i\}_{i \in I} \vdash^n xp\{l_i : P_i\}_{i \in I}} \quad [\text{T-LABEL}]$	$\frac{\Gamma \vdash^{n_1} P \quad \Gamma \vdash^{n_2} Q}{\Gamma \vdash^{n_k} P \oplus_k Q} \quad [\text{T-CHOICE}]$	
$\frac{\Gamma, x : T \vdash^n P}{\Gamma, x : S \vdash^{1+n} [x]P} \quad [\text{T-CAST}]$	$\frac{\Gamma, x : S_k \vdash^n P_k}{\Gamma, x : p\{l_i : S_i\}_{i \in I} \vdash^n xp\{l_i : P_i\}_{i \in I}} \quad [\text{CO-LABEL}] \quad k \in I$	$\frac{\Gamma \vdash^n P_k}{\Gamma \vdash^n P_1 \oplus_k P_2} \quad [\text{CO-CHOICE}]$	
$\frac{\Gamma, x : S \vdash^m P \quad \Delta, x : T \vdash^n Q}{\Gamma, \Delta \vdash^{1+m+n} (x)(P \mid Q)} \quad [\text{T-PAR}] \quad S \sim T$	$\frac{\overline{x : \bar{S} \vdash^n P}}{x : \bar{S} \vdash^{m+n} A\langle \bar{x} \rangle} \quad [\text{T-CALL}] \quad A : [\bar{S}; n], A(\bar{x}) \triangleq P$		

6 THE TYPE SYSTEM, FORMALLY

The typing rules resemble those of a traditional session type system but differ in a few key aspects. First of all, they establish a tighter-than-usual correspondence between types and processes so that any discrepancy between actual and expected types is accounted for by explicit casts. This way, we make sure that actions leading to the termination of a session *at the type level* are matched by corresponding actions *at the process level*, a key property used in the soundness proof of the type system (Theorem 6.4). In addition, the typing rules enforce the boundedness properties informally described in the previous section. Action boundedness is enforced by specifying the typing rules as a generalized inference system and using two corules to make sure that every well-typed process is at finite distance from **done** or a **close** x . Concerning session and cast boundedness, we annotate typing judgments with a *rank*, that is an upper bound to the number of casts that must be performed and of sessions that must be created in order to terminate the process in the judgment.

The typing rules are defined by the generalized inference system in Table 3 and derive judgements of the form $\Gamma \vdash^n P$, meaning that P is well typed in the *typing context* Γ and has rank n . A typing context is a finite map from channels to session types written $x_1 : S_1, \dots, x_n : S_n$ or $x : \bar{S}$. We use Γ and Δ to range over typing contexts, we write $\emptyset$ for the empty context and Γ, Δ for the union of Γ and Δ when they have disjoint domains. We type check a program $\{A_i(\bar{x}_i) \triangleq P_i\}_{i \in I}$ under a global set of type assignments $\{A_i : [\bar{S}_i; n_i]\}_{i \in I}$ associating each process name A_i with a tuple of session types $\bar{S}_i$ and a rank n_i . The program is well typed if $\bar{x}_i : \bar{S}_i \vdash^{n_i} P_i$ for every $i \in I$, establishing that the tuple $\bar{S}_i$ corresponds to the way the channels $\bar{x}_i$ are used by P_i and that n_i is a feasible rank annotation for P_i . Hereafter, we omit the rank from judgments when it is not important.

Let us look at the typing (co)rules in detail. **[T-DONE]** is the usual axiom requiring that the terminated process leaves no unused channels behind. Since **done** performs no casts and creates no sessions, it can have any rank. Rules **[T-WAIT]** and **[T-CLOSE]** concern the exchange of session termination signals. There is nothing remarkable here except noting once again that the rank of **close** x can be arbitrary. Rules **[T-CHANNEL-IN]** and **[T-CHANNEL-OUT]** are similar, but they concern the exchange of channels. Note that, in **[T-CHANNEL-OUT]**, the type T of the message y is required to match *exactly* that in the type of the channel x used for the communication, whereas **Gay and Hole**

[2005] allow the type of y to be a subtype of T . This is one instance of the “tight correspondence” that we mentioned earlier. The rule [T-LABEL] deals with the input/output of labels. As usual, any channel other than the one affected by the communication must be used in exactly the same way in every branch. However, the rule is stricter than that of Gay and Hole [2005] because it requires an exact correspondence between the labels that can be exchanged on x by the process and those in the type of x . The fact that a conclusion and premises are all annotated with the same rank n means that n is an upper bound for the rank of all branches of a label input/output. The corule [CO-LABEL] does not impose additional constraints compared to [T-LABEL] and has *exactly one premise*, corresponding to one branch of the process in the conclusion. The effect of [CO-LABEL], when interpreted inductively together with the other rules, is to ensure the existence of a finite typing derivation whose leaves are applications of [T-DONE] or [T-CLOSE], hence action boundedness.

Rule [T-CHOICE] is a standard typing rule for non-deterministic choices, requiring that both branches are well typed in exactly the same typing context. Notice that the rank of a choice $P_1 \oplus_k P_2$ is determined by the branch indexed by the k annotation, which is elected as the branch that leads to termination (see also Remark 4 for a comparison with [T-LABEL]). Like [CO-LABEL], the associated corule [CO-CHOICE] ensures that the same branch gets closer to *done* or a *close* x to enforce action boundedness. Without this corule, it would not be possible to find a *finite-depth* derivation tree for an action-bounded process such as A in Eq. (3). Coherently with [T-CHOICE], the same branch that leads to termination is also the one that determines the rank of the choice as a whole.

Rule [T-CAST] is Liskov’s substitution principle formulated as an inference rule. It states that a channel x of type S can be safely used where a channel of type T is expected, provided that $S \leq T$. The most important detail to notice here is that the rank of a cast is one plus that of the process in which the cast has effect. This way we account for this cast in the rank of the process so as to guarantee cast boundedness. Rule [T-PAR] concerns parallel composition and session creation. The rule is shaped after the cut rule of linear logic also adopted in other session type systems based on linear logic [Caires et al. 2016; Lindley and Morris 2016; Wadler 2014]. In particular, the parallel processes P and Q share no channel other than the session x that connects them, so as to prevent mutual dependencies between sessions and guarantee deadlock freedom. The side condition $S \sim T$ requires that the way in which P and Q use channel x is such that the session x can fairly terminate (cf. Definition 3.7). We *do not* require that S and T are dual to each other because reductions (cf. [R-PICK]) and structural pre-congruence (cf. [S-CAST-NEW]) do not necessarily preserve session type duality. Also, duality does not always imply compatibility (Section 3.5). The rank of a parallel composition is one plus that of the composed processes. By accounting for each occurrence of parallel compositions in the rank, we guarantee that well-typed processes are session bounded.

Finally, rule [T-CALL] states that a process invocation $A(\bar{x})$ is well typed provided that the types associated with $\bar{x}$ match those of the global assignment $A : [\bar{S}; n]$. Note that [T-CALL] is *not* an axiom: its premise (re)checks that the body P in the definition of A is coherent with the global type assignment $A : [\bar{S}; n]$. With this formulation of [T-CALL], the only axioms are [T-DONE] and [T-CLOSE] so that the inductive interpretation of the typing (co)rules ensures action boundedness. Note also that the rank of the conclusion may be greater than the rank n associated with A . This overapproximation grants more flexibility when typing different branches in [T-LABEL].

Example 6.1. To show that the program defined in Example 4.1 is well typed, consider the session types $S = ?\text{add}.S + ?\text{pay}.\text{end}$ and $T = !\text{add}.(!\text{add}.T \oplus !\text{pay}.\text{end})$ and the global type assignments $A : [T; 0]$, $B : [S, !\text{ship}.\text{end}; 0]$, $C : [?\text{ship}.\text{end}; 0]$ and $\text{Main} : [(); 3]$. We can obtain typing

for the definition of A and

$$\frac{\vdots \quad \frac{\overline{x : T \vdash^0 A \langle x \rangle} \text{ [T-CALL]} \quad \overline{x : !\text{end} \vdash^0 \text{close } x} \text{ [T-CLOSE]}}{x : !\text{add}.T \oplus !\text{pay}.\text{!end} \vdash^0 x!\{\text{add} : A \langle x \rangle, \text{pay} : \text{close } x\}} \text{ [T-LABEL]} \quad \frac{}{x : T \vdash^0 x!\text{add}.x!\{\text{add} : A \langle x \rangle, \text{pay} : \text{close } x\}} \text{ [T-LABEL]}$$

for the definition of A and

$$\frac{\displaystyle \frac{\displaystyle \vdots}{x : S, y : !\mathsf{ship}.\mathsf{end} \vdash^0 B\langle x, y \rangle} [\text{T-CALL}] \quad \frac{\displaystyle \frac{\displaystyle \frac{\displaystyle y : !\mathsf{end} \vdash^0 \mathsf{close } y}{} [\text{T-CLOSE}] \quad \frac{\displaystyle y : !\mathsf{ship}.\mathsf{end} \vdash^0 y!\mathsf{ship}.\mathsf{close } y}{} [\text{T-LABEL}]}{\displaystyle y : !\mathsf{end}, y : !\mathsf{ship}.\mathsf{end} \vdash^0 \mathsf{wait } x.y!\mathsf{ship}.\mathsf{close } y} [\text{T-WAIT}]} [\text{T-LABEL}]}{x : S, y : !\mathsf{ship}.\mathsf{end} \vdash^0 x?\{\mathsf{add} : B\langle x, y \rangle, \mathsf{pay} : \mathsf{wait } x.y!\mathsf{ship}.\mathsf{close } y\}}$$

$$\frac{\frac{\vdots}{x : T \vdash^0 A\langle x \rangle} \text{[T-CALL]} \quad \frac{\vdots}{x : S^\perp \vdash^1 [\bar{x}]A\langle x \rangle} \text{[T-CAST]} \quad \frac{\vdots}{x : S, y : !\text{ship}!\text{end} \vdash^0 B\langle x, y \rangle} \text{[T-CALL]} \quad \vdots}{y : !\text{ship}!\text{end} \vdash^2 (x)([\bar{x}]A\langle x \rangle \mid B\langle x, y \rangle) \quad y : ?\text{ship}?\text{end} \vdash^0 C\langle y \rangle} \text{[T-PAR]} \\ \frac{}{\emptyset \vdash^3 (y)((x)([\bar{x}]A\langle x \rangle \mid B\langle x, y \rangle) \mid C\langle y \rangle)}$$

showing that *Main* too is well typed. In all cases, we have truncated the proof trees above the applications of [T-CALL]. Of course, for each judgment occurring in these proof trees, we also have to exhibit a finite proof tree possibly using [CO-LABEL] proving action boundedness. This can be easily achieved for the given process definitions, observing that none of *A*, *B* and *C* creates new sessions and that all of their typing derivations have a finite branch. $\lrcorner$

Example 6.2. In this example we demonstrate that well-typed processes may still create an unbounded number of (nested) sessions. To this aim, let us consider again the process C defined in Eq. (6). Notice that C is a choice whose left branch creates a new session and whose right branch does not. For this reason, we elect the right choice as the one that leads to termination, and therefore that determines the rank of the process. We derive

$$\frac{\frac{\vdots}{y : !\text{end} \vdash^0 C\langle y \rangle} [\text{T-CALL}] \quad \frac{\frac{x : !\text{end} \vdash^0 \text{close } x}{} [\text{T-CLOSE}] \quad \frac{}{x : !\text{end}, y : ?\text{end} \vdash^0 \text{wait } y.\text{close } x} [\text{T-WAIT}]}{x : !\text{end} \vdash^1 (y)(C\langle y \rangle \mid \text{wait } y.\text{close } x)} [\text{T-PAR}] \quad \frac{}{x : !\text{end} \vdash^0 \text{close } x}}{x : !\text{end} \vdash^0 (y)(C\langle y \rangle \mid \text{wait } y.\text{close } x) \oplus_2 \text{close } x}$$

In a similar way, there exist well-typed processes that perform an unbounded number of casts but whose rank is finite. For example, it is easy to obtain a typing derivation for the following alternative definition of the process A discussed in Example 4.1:

$$A(x) \triangleq x!\text{add}.(\llbracket x \rrbracket x!\text{add}.A\langle x \rangle \oplus_2 \llbracket x \rrbracket x!\text{pay.close } x)$$

Even though this process uses fair subtyping an unbounded number of times, the right branch of the choice has rank 1, which is all we need to conclude that the process has rank 1 overall. $\square$

Example 6.3 (random bit generator). Below we define a process $A(x)$ that implements the random bit generator protocol discussed in Examples 3.5 and 3.11 along with a client $B(x)$ that asks the service for random bits until it receives a 1:

$$\begin{aligned} A(x) &\triangleq x!\{\text{more} : x!\{0 : A\langle x \rangle, 1 : A\langle x \rangle\}, \text{stop} : \text{close } x\} \\ B(x) &\triangleq x!\text{more}.x?\{0 : B\langle x \rangle, 1 : x!\text{stop}.wait\ x.\text{done}\} \end{aligned}$$

These definitions are well typed under the global assignments $A : [S; 0]$ and $B : [U; 0]$ where S is defined as in Example 3.5 and U is defined as in Example 3.11. Notice that the termination of the parallel composition of A and B depends on a complex negotiation between A and B . Indeed, A terminates when it receives 0 from B and B terminates when it receives stop from A . This interaction pattern can only be modeled if both A and B are defined using general recursion. $\dashv$

Well-typed processes enjoy the expected properties, including typing preservation under structural pre-congruence and reduction. Most importantly, they fairly terminate:

THEOREM 6.4. *If $\emptyset \vdash P$, then P is fairly terminating.*

The most intriguing aspect in the proof of Theorem 6.4 is that a closed, well-typed process admits a reduction sequence to **done**. Space constraints force us to relegate the details to the Appendix in the supplemental material, so here we only sketch the key elements of the proof, which is related to the *method of helpful directions* [Francez 1986]: we define a well-founded *measure* for (well-typed) processes and we prove that this measure decreases strictly as the result of “helpful” reductions. In our case, the measure of a (well-typed) process P is a lexicographically ordered pair (m, n) of natural numbers such that m is an upper bound to the number of sessions that P may need to create and of casts that P may need to perform *in the future* in order to terminate, whereas n is the cumulative rank of the sessions that P has created *in the past* and that are not terminated yet. A session terminates by **[R-CLOSE]**; a cast is performed when it is absorbed by the corresponding restriction, namely by **[S-CAST-NEW]**. To distinguish between past and future of P we look at its structure: all sessions that occur unguarded in P have been created and are not terminated; all casts that occur in P are yet to be performed; all sessions that occur guarded in P have not been created yet. To compute the measure of a process, we introduce three refined typing rules to derive judgments of the form $\Gamma \models^\mu P$, stating that P is well typed in Γ and has measure μ :

$$\begin{array}{c} \text{[MT-THREAD]} \\ \frac{\Gamma \vdash^n P}{\Gamma \models^{(n,0)} P} \end{array} \quad \begin{array}{c} \text{[MT-PAR]} \\ \frac{\Gamma, x : S \models^\mu P \quad \Delta, x : T \models^\nu P}{\Gamma, \Delta \models^{\mu+\nu+(0, \|S, T\|)} (x)(P \mid Q)} S \sim T \end{array} \quad \begin{array}{c} \text{[MT-CAST]} \\ \frac{\Gamma, x : T \models^\mu P}{\Gamma, x : S \models^{\mu+(1,0)} [x]P} S \leq T \end{array}$$

Rule **[MT-THREAD]** has *lower priority* than **[MT-PAR]** and **[MT-CAST]**, in the sense that it applies only to processes that are not a cast or a parallel composition. We call such processes *threads* and their measure is solely determined by their rank: every cast occurring in a thread is yet to be performed and every session occurring in a thread is yet to be created. Rule **[MT-PAR]** states that the measure of a parallel composition is the (pointwise) sum of the measures of the composed processes, taking into account the rank of the session x by which they are connected. Finally, **[MT-CAST]** states that the measure of a cast is the same measure of the process in which the cast has effect, but with the first component increased by one to account for the fact that the cast is yet to be performed.

Note that, as a well-typed process reduces, its measure may vary arbitrarily. In particular, its measure *may increase* if the process chooses to create new sessions (*cf.* the left choice of process C in Eq. (6)) or if it picks a label that lengthens the shortest path leading to session termination (*cf.* Example 4.3). Nonetheless, the key lemma below assures that the measure of a well-typed process *may also decrease* following carefully chosen, “helpful” reductions.

LEMMA 6.5. *If $\emptyset \models^\mu P$, then either $P \leq \text{done}$ or $P \Longrightarrow^+ Q$ and $\emptyset \models^\nu Q$ for some Q and $\nu < \mu$.*

Remark 4. The rank of a non-deterministic choice $P \oplus Q$ can usually be chosen to be the minimum among those of the branches P and Q , so that the type system can handle processes like those in Example 6.2, which *may* create new sessions or perform casts but they need not do so in order to terminate. On the contrary, the rank of a label output $x!\{l_i : P_i\}_{i \in I}$ has to be an upper bound of that of all branches P_i . The motivation for such different ways of determining the rank of these process forms, despite both represent an *internal choice*, lies in the proof of Lemma 6.5. In $P \oplus Q$, both branches are typed in *exactly the same* typing context, meaning that the choice of one branch or the other has no substantial impact on the shortest paths that terminate the sessions used by P and Q . Thus, the “helpful” reduction can be solely driven by the rank of the chosen branch. In a label output $x!\{l_i : P_i\}_{i \in I}$ it could happen that all branches with minimum rank increase the length of the shortest path that leads to the termination of x . In this case, the choice of the “helpful” reduction must prioritize the termination of x , but then the rank of the whole process has to be an upper bound of that of the branches to be sure that the measure of the reduct decreases. $\square$

7 ON FAIR SUBTYPING AND HIGHER-ORDER SESSION TYPES

We have defined fair subtyping is such a way that higher-order session types are *invariant* with respect to the type of the channel being exchanged (cf. [F-CHANNEL]). This is a limitation compared to traditional presentations of unfair subtyping [Bernardi and Hennessy 2016; Castagna et al. 2009; Gay and Hole 2005], where the covariant/contravariant rules shown below are adopted:

$$\frac{[U\text{-CHANNEL-IN}]}{U \leq V \quad S \leq T} \quad \frac{[U\text{-CHANNEL-OUT}]}{V \leq U \quad S \leq T} \\ \frac{}{?U.S \leq ?V.T} \quad \frac{}{!U.S \leq !V.T}$$

The problem of these rules is that a single application of fair subtyping allowing for co-/contravariance of higher-order session types may have the same overall effect of infinitely many applications of fair subtyping on first-order session types and, as we have seen in Section 5.3, unbounded applications of fair subtyping may compromise fair termination. Below is an example that illustrates the problem. The example is not large *per se*, but it is a bit contrived because it has to involve two sessions (or else there would be no need for higher-order session types), it must be bounded (or else it could be ruled out by the action/session/cast boundedness requirements) and non-terminating:

$$(y)((x)(A\langle x, y \rangle \mid B\langle x \rangle) \mid B\langle y \rangle) \text{ where } \begin{array}{l} A\langle x, y \rangle \triangleq x!\text{more}.x!y.B\langle x \rangle \\ B\langle x \rangle \triangleq x?\{\text{more} : x?(y).A\langle y, x \rangle, \text{stop} : \text{wait } x.\text{done}\} \end{array} \quad (10)$$

The process models a *master* $A\langle x, y \rangle$ connected with a *primary slave* $B\langle x \rangle$ and a *secondary slave* $B\langle y \rangle$ through the sessions x and y . The interaction among the three processes proceeds in rounds. At each round, the master may decide whether to continue or stop the interaction by sending either *more* or *stop* on the session x to the primary slave. If the master decides to continue the interaction (which it does deterministically), it also delegates y to the primary slave so that, at the next round, the roles of the three processes rotate: the master is downgraded to secondary slave, the primary slave is promoted to master, and the secondary slave becomes the primary one. Below is a graphical representation of the network topology modeled by the process and of its evolution:

It is clear that the process in Eq. (10) does not terminate since there is no *close* x to match the *wait* x . It is also relatively easy to infer the types of x and y from the structure of $A\langle x, y \rangle$ and $B\langle x \rangle$.

In particular, if we call S_A and T_A the types of x and y in $A(x, y)$ and S_B the type of x in $B(x)$ we see that these types must satisfy the equations

$$S_A = !\text{more}.\text{!}T_A.S_B \quad S_B = ?\text{more}.\text{?}S_A.T_A + ?\text{stop}.\text{?end} \quad T_A = !\text{more}.\text{!}T_A.S_B \oplus !\text{stop}.\text{!end}$$

Note that $T_A \leq S_A$ holds because T_A and S_A differ only for the topmost output. The validity of this relation is unquestionable as it relies on the definition of fair subtyping that we have given in Section 3.2, which is invariant with respect to higher-order session types. If fair subtyping allowed for covariance of higher-order inputs (cf. [U-CHANNEL-IN]), then $T_A^\perp \leq S_B$ would also hold and we would be able to establish that the process in Eq. (10) is well typed, provided that casts are placed appropriately. Below we show a version of the process in which we have annotated restrictions with the types $S \mid T$ of the two endpoints and casts with the target type of the channel affected by subtyping. The interested reader can find a full typing derivation in the supplemental material:

$$(y : T_A \mid T_A^\perp)((x : T_A \mid T_A^\perp)([x : S_A]A\langle x, y \rangle \mid [x : S_B]B\langle x \rangle) \mid [y : S_B]B\langle y \rangle) \quad (11)$$

Dually, if fair subtyping allowed for contravariance of higher-order outputs (cf. [U-CHANNEL-OUT]) then $S_B^\perp \leq T_A$ would also hold (along with $S_B^\perp \leq S_A$ by transitivity of $\leq$) and we would be able to establish that the above process is well typed according to the type annotations shown below:

$$(y : S_B^\perp \mid S_B)((x : S_B^\perp \mid S_B)([x : S_A][y : T_A]A\langle x, y \rangle \mid B\langle x \rangle) \mid B\langle y \rangle)$$

By restricting fair subtyping of higher-order session types to invariant inputs and outputs, the only chance we have to build a typing derivation for the process in Eq. (10) is by casting y each time it is delegated, either before it is sent or after it is received, for example as in

$$\begin{aligned} A(x : U_A, y : V_A) &\triangleq x!\text{more}.\text{!}[y : U_A]x!y.B\langle x \rangle \\ B(x : U_B) &\triangleq x?\{\text{more} : x?(y : U_A).A\langle y, x \rangle, \text{stop} : \text{wait } x.\text{done}\} \end{aligned}$$

where

$$U_A = !\text{more}.\text{!}U_A.U_B \quad U_B = ?\text{more}.\text{?}U_A.V_A + ?\text{stop}.\text{?end} \quad V_A = !\text{more}.\text{!}U_A.U_B \oplus !\text{stop}.\text{!end}$$

Note that $[y : U_A]$ is a “first-order” cast, in the sense that the relation $V_A \leq U_A$ holds for fair subtyping as defined in Section 3.2 without using [U-CHANNEL-IN] or [U-CHANNEL-OUT], but the cast is now placed in a region within the definition of A that prevents finding a finite rank for A .

8 RELATED WORK

Deadlock freedom. The absence of deadlocks is a fundamental requirement for the proof of Theorem 6.4. In this work we ensure deadlock freedom by adopting the formulation of rule [T-PAR] inspired to linear logic, as in the works of Caires et al. [2016], Wadler [2014] and Lindley and Morris [2016]. This technique is simple and effective, but limits its applicability to tree-shaped network topologies. Another line of works makes use of a richer type structure associating input/output actions in types with “levels” or “priorities” so as to prevent circular dependencies. This approach has been pioneered by Kobayashi [2002, 2006] in the π -calculus and then ported to session-based calculi by Padovani [2014] and Dardha and Gay [2018]. Dardha and Pérez [2015] compare the two approaches. Balzer et al. [2019] relax the approach based on linear logic so as to allow controlled forms of channel sharing but still preserving deadlock freedom. Kobayashi and Laneve [2017] describe another approach based on behavioral types associated to processes (as opposed to channels) that is capable of addressing cyclic network topologies. Interestingly, the proof of Theorem 6.4 is independent of the technique used for ensuring deadlock freedom, hence our type system can be combined with other approaches addressing more complex network topologies.

Lock freedom. Type systems ensuring this liveness property have been studied by Kobayashi [2002] and Kobayashi and Sangiorgi [2010] for the π -calculus, by Padovani [2014] for the linear

π -calculus and calculi based on binary sessions, and by Padovani et al. [2014] for the conversation calculus, a calculus of multiparty sessions. A common trait of the type systems in these works, which is also the main distinguishing aspect with our own, is that they use explicit type annotations to capture the dependencies between different actions. In particular, the annotations in the works of Kobayashi [2002], Kobayashi and Sangiorgi [2010] and Padovani [2014] conservatively estimate the amount of reductions that are necessary in order to perform a given input/output action. If such amounts cannot be statically determined or if they are not finite, as is the case of the action $y!\text{ship}$ in Eq. (1) and in general in all (recursive) processes where an action may be preceded by an arbitrary number of other actions, either the process is declared ill typed or the action is not guaranteed to be eventually performed. As a matter of fact, determining the amount of reductions that are necessary to unlock and perform a certain action is crucially important also in our work, but this information stays buried in the soundness proof of the type system as part of the measure and does not surface in types. The use of unannotated types allows our type system to address a broad family of processes – of which Eq. (1) is just one representative member – for which the existing type-based techniques are unable to guarantee lock freedom.

Termination by typing. Session type systems strictly based on linear logic [Caires et al. 2016; Wadler 2014] ensure session termination, but they do not allow for the specification of recursive protocols. Yoshida et al. [2004] present a type system for the π -calculus based on linear channel types that ensures strong normalization, hence lock freedom. Their type system applies also to binary sessions given the tight relationship between binary sessions and linear channel types [Dardha et al. 2017]. Deng and Sangiorgi [2006] and Demangeon et al. [2009] study type-based approaches for enforcing the termination of (mobile) processes with replication, namely the property that every reduction sequence is guaranteed to be finite. These type systems could be applied also to session-based calculi. Kobayashi and Sangiorgi [2010] show that the combination of deadlock freedom and termination results in lock freedom. Lindley and Morris [2016] study languages of binary sessions for which typing guarantees strong normalization, implying that well-typed programs are also lock free. Their type system is based on an extension of linear logic with fixed points [Baelde 2012; Doumane 2017] so that recursive types come into dual forms corresponding to least and greatest fixed points. Programs like those discussed in Examples 4.1 and 6.3, which fairly terminate but admit infinite computations, would be ill typed if modeled in their languages.

Multiparty sessions. Multiparty sessions [Honda et al. 2008, 2016] are a generalization of binary sessions to an arbitrary – sometimes variable – number of participants. In recent advancements to the theory of multiparty session types, Scalas and Yoshida [2019] show how to specify and enforce a family of safety and liveness properties of multiparty sessions as formulas in the modal μ -calculus. In particular, they define three predicates *live*, *live*⁺ and *live*⁺⁺ closely related to lock freedom: while *live*⁺ and *live*⁺⁺ specify “bounded” forms of lock freedom such that every pending communication is guaranteed to be performed in a finite number of steps, *live* specifies the eventual completion of *any* pending communication action. If we were to reason about process *B* in Eq. (1) using these predicates, we could say that *B* satisfies *live* but not *live*⁺ or *live*⁺⁺. Their type system is able to enforce *live*⁺ and *live*⁺⁺, but not *live* because they use an “unfair” subtyping relation that does not preserve (all) liveness properties. More recently, van Glabbeek et al. [2021] have proposed a type system for multiparty sessions that ensures lock freedom and that is not only sound (under suitable conditions) but also *complete*, in the sense that all lock-free sessions are well typed. Such a strong result is possible because the notion of lock freedom they consider is based on *justness* [van Glabbeek 2019; van Glabbeek and Höfner 2019], a “minimal fairness assumption that guarantees only that concurrent transitions cannot prevent each other from happening” [van Glabbeek et al. 2021]. We assume strong fairness, which is at the opposite end of the spectrum compared to justness. The stronger the fairness assumption, the larger the set of “unfair” executions that are ruled out, the

larger the family of lock-free processes. For example, the process modeled in Eq. (1) is not lock free under justness. Another difference between our work and the one of [van Glabbeek et al. \[2021\]](#) is that they only consider single, first-order session systems (even though sessions can be multiparty). This simplification avoids all the issues arising from session creation and delegation (Sections 5 and 7) and from the interleaving of blocking actions from different sessions, which occurs in Eq. (1) and is one of the motivations for focusing on fair termination instead of lock freedom.

Fair subtyping. The original “unfair” subtyping relation for session types of [Gay and Hole \[2005\]](#) preserves communication safety but not necessarily (fair) termination. Analogous relations have been studied by [Castagna et al. \[2009\]](#) and [Bernardi and Hennessey \[2016\]](#). Fair subtyping is a liveness-preserving refinement of unfair subtyping. [Bravetti and Zavattaro \[2009\]](#) study a *strong subcontract relation* for Web service contracts that shares the same liveness-preserving features of fair subtyping. Indeed, the notion of *strong contract composition* of [Bravetti and Zavattaro \[2009\]](#) that is preserved by strong subcontract is a first-order version of compatibility (Definition 3.7) for multiparty service compositions. Both strong subcontract and fair subtyping are closely related to divergence-insensitive refinements for processes called *fair testing* [[Natarajan and Cleaveland 1995](#)] and *should testing* [[Rensink and Vogler 2007](#)]. Inference systems for fair subtyping have been studied by [Padovani \[2013, 2016\]](#) and [Ciccone and Padovani \[2021b\]](#). The generalized inference system for fair subtyping in Table 1 is a higher-order adaptation of the one presented by [Ciccone and Padovani \[2021b\]](#). [Bravetti et al. \[2021\]](#) study a fair subtyping relation for *asynchronous* session types (in which output actions are non-blocking), showing that it is undecidable and proposing decidable approximations of it. It appears feasible that asynchronous fair subtyping could be used in our type system as a drop-in replacement of fair subtyping to further relax the correspondence between the structure of protocols and processes.

Fair termination. [Grumberg et al. \[1984\]](#) describe proof methods for proving the fair termination of communicating processes specified as CSP terms, hence for a process model comparable to that of our (first-order) session types. They point out how fair termination enables the characterization of behaviors like that of Eq. (1), where processes engaged in a finite but unbounded number of communications may (fairly) terminate. [Cook et al. \[2007\]](#) check liveness properties of C programs by means of a reduction to a fair termination problem, whereas [Ganty et al. \[2009\]](#) give a procedure for deciding fair termination of event-driven programs. [Tassarotti et al. \[2017\]](#) present an extension of separation logic to prove the correctness of an implementation of binary sessions. Interestingly, they also rely on a liveness-preserving refinement, although in their case the refinement is used to compare specification and implementation of the same program rather than protocols. The work of [Cacciagrano et al. \[2006, 2009\]](#) shows that a notion of fair process termination for the π -calculus formulated in terms of fair runs may differ from the one we have adopted in Definition 4.2. Given that the counterexample they use relies crucially on replication and shared channels, it might be interesting to investigate whether it is possible to restore the equivalence of the two formulations for our (typed) language, which is based on recursion and linearized channels.

9 CONCLUDING REMARKS

We have presented the first type system ensuring the fair termination of binary sessions (Theorem 6.4). The type system applies to a calculus that supports general recursion, session interleaving, session delegation and dynamic session creation, thus targeting a large family of processes with unboundedly many states and communicating in networks with variable topology. Fair termination is coarser than strong normalization or plain termination, since it does not rule out infinite interactions, and it is stricter than lock freedom, so it entails the eventual completion of *any* pending communication, including those blocked by an unbounded number of actions in possibly different (Eq. (1)) or yet-to-be-created (Eq. (6)) sessions. These interaction patterns fall outside the scope of

existing type systems for lock freedom. Clearly, there exist lock-free processes that are not fairly terminating (Section 1), although the possibility that interactions may eventually terminate seems to be a natural requirement for sessions. For all of these reasons, we consider our type system a substantial leap forward in the study of type-based techniques for the enforcement of liveness properties of sessions.

A key element of the type system is fair subtyping, a liveness-preserving subtyping relation for session types that so far has only been studied theoretically but never applied in a type system. We have shown that fair subtyping must be used with care (Section 5.3). Indeed, our type system accounts for all the usages of fair subtyping, making sure that the overall effort required to terminate a process, along with all the sessions it has created, remains finite. We have also uncovered an unforeseen interaction between fair subtyping and higher-order sessions that may invalidate the liveness-preserving features of fair subtyping (Section 7). One simple way to avoid this issue is by imposing the invariance of higher-order session types (Section 3.2). While this is a restriction compared to the “unfair” subtyping relations for session types, we think it is unlikely to have a profound impact: higher-order channels in distributed programs are far less ubiquitous than, say, higher-order functions in functional ones, and co-/contra-variance of higher-order session types can be partly recovered by means of explicit casts, at least in those places where casts can be used.

We could not detail any algorithmic aspect of the type system in the main body of the paper because of space limits. In particular, the formulation of the typing rules in Table 3, in which the rank annotation is essentially guessed, allows for simple presentation and soundness proof of the type system, but does not say how to determine the rank annotation nor when it exists. We have defined an algorithm for computing the *minimum rank annotation* that is needed to type a process as well as necessary and sufficient conditions for its existence that are solely based on the structure of processes. This makes it possible to provide an equivalent and fully compositional set of typing rules with which we have proved that type checking is decidable if one assumes that bindings and casts are decorated with explicit type annotations. A proof-of-concept Haskell implementation of this type checking algorithm is available [Ciccone and Padovani 2021a]. We have also proved that the location of casts can be inferred automatically and have defined a sound and complete set of *co-contextual typing rules* [Erdweg et al. 2015] to reconstruct session types from unannotated processes. Details about these aspects can be found in the supplemental material associated with this paper.

Considering that fair subtyping for binary and multiparty session types share essentially the same characterization [Bravetti and Zavattaro 2009; Bugliesi et al. 2009; Padovani 2016], we expect the type system to scale smoothly to multiparty sessions [Honda et al. 2008, 2016]. It remains to be established if fair subtyping can also be applied in the general framework proposed by Scalas and Yoshida [2019], which is parametric in the liveness property one is interested to enforce, or if different versions of fair subtyping are necessary to enforce different liveness properties. As suggested by one reviewer, another interesting direction for future investigations is the application of the type system to higher-order session calculi such as GV [Gay and Vasconcelos 2010; Wadler 2014] and μ GV [Lindley and Morris 2016]. In particular, establishing the rank of terms in a higher-order calculus appears to be less straightforward than in the process calculus that we consider.

ACKNOWLEDGMENTS

We are grateful to the anonymous POPL reviewers for their thoughtful comments and suggestions. A special mention goes to the members of the Artifact Evaluation Committee who evaluated our artifact. The extent and quality of their feedback has been invaluable to improve the usability and the documentation of the artifact.

REFERENCES

- Peter Aczel. 1977. An Introduction to Inductive Definitions. In *Handbook of Mathematical Logic*, Jon Barwise (Ed.). Studies in Logic and the Foundations of Mathematics, Vol. 90. Elsevier, 739 – 782. [https://doi.org/10.1016/S0049-237X\(08\)71120-0](https://doi.org/10.1016/S0049-237X(08)71120-0)
- Davide Ancona, Francesco Dagnino, and Elena Zucca. 2017. Generalizing Inference Systems by Coaxioms. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10201)*, Hongseok Yang (Ed.). Springer, 29–55. https://doi.org/10.1007/978-3-662-54434-1_2
- Krzysztof R. Apt, Nissim Francez, and Shmuel Katz. 1987. Appraising Fairness in Languages for Distributed Programming. In *Conference Record of the Fourteenth Annual ACM Symposium on Principles of Programming Languages, Munich, Germany, January 21-23, 1987*. ACM Press, 189–198. <https://doi.org/10.1145/41625.41642>
- David Baelde. 2012. Least and Greatest Fixed Points in Linear Logic. *ACM Trans. Comput. Log.* 13, 1 (2012), 2:1–2:44. <https://doi.org/10.1145/2071368.2071370>
- Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. 2019. Manifest Deadlock-Freedom for Shared Session Types. In *Programming Languages and Systems - 28th European Symposium on Programming, ESOP 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11423)*, Luís Caires (Ed.). Springer, 611–639. https://doi.org/10.1007/978-3-030-17184-1_22
- Giovanni Bernardi and Matthew Hennessy. 2016. Using higher-order contracts to model session types. *Log. Methods Comput. Sci.* 12, 2 (2016). [https://doi.org/10.2168/LMCS-12\(2:10\)2016](https://doi.org/10.2168/LMCS-12(2:10)2016)
- Philippe Bidinger and Adriana B. Compagnoni. 2009. Pict correctness revisited. *Theor. Comput. Sci.* 410, 2-3 (2009), 114–127. <https://doi.org/10.1016/j.tcs.2008.09.014>
- Mario Bravetti, Julien Lange, and Gianluigi Zavattaro. 2021. Fair Refinement for Asynchronous Session Types. In *Foundations of Software Science and Computation Structures - 24th International Conference, FOSSACS 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 (Lecture Notes in Computer Science, Vol. 12650)*, Stefan Kiefer and Christine Tasson (Eds.). Springer, 144–163. https://doi.org/10.1007/978-3-030-71995-1_8
- Mario Bravetti and Gianluigi Zavattaro. 2009. A theory of contracts for strong service compliance. *Math. Struct. Comput. Sci.* 19, 3 (2009), 601–638. <https://doi.org/10.1017/S0960129509007658>
- Michele Bugliesi, Damiano Macedonio, Luca Pino, and Sabina Rossi. 2009. Compliance Preorders for Web Services. In *Web Services and Formal Methods, 6th International Workshop, WS-FM 2009, Bologna, Italy, September 4-5, 2009, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 6194)*, Cosimo Laneve and Jianwen Su (Eds.). Springer, 76–91. https://doi.org/10.1007/978-3-642-14458-5_5
- Diletta Cacciagrano, Flavio Corradini, and Catuscia Palamidessi. 2006. Fair Pi. In *Proceedings of the 13th International Workshop on Expressiveness in Concurrency, EXPRESS 2006, Bonn, Germany, August 26, 2006 (Electronic Notes in Theoretical Computer Science, Vol. 175)*, Roberto M. Amadio and Iain Phillips (Eds.). Elsevier, 3–26. Issue 3. <https://doi.org/10.1016/j.entcs.2006.10.051>
- Diletta Cacciagrano, Flavio Corradini, and Catuscia Palamidessi. 2009. Explicit fairness in testing semantics. *Log. Methods Comput. Sci.* 5, 2 (2009). [https://doi.org/10.2168/LMCS-5\(2:15\)2009](https://doi.org/10.2168/LMCS-5(2:15)2009)
- Luís Caires, Frank Pfenning, and Bernardo Toninho. 2016. Linear logic propositions as session types. *Math. Struct. Comput. Sci.* 26, 3 (2016), 367–423. <https://doi.org/10.1017/S0960129514000218>
- Giuseppe Castagna, Mariangiola Dezani-Ciancaglini, Elena Giachino, and Luca Padovani. 2009. Foundations of session types. In *Proceedings of the 11th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, September 7-9, 2009, Coimbra, Portugal*, António Porto and Francisco Javier López-Fraguas (Eds.). ACM, 219–230. <https://doi.org/10.1145/1599410.1599437>
- Luca Ciccone and Luca Padovani. 2021a. *FairCheck*. Retrieved October 1, 2021 from <https://github.com/boystrange/FairCheck>
- Luca Ciccone and Luca Padovani. 2021b. Inference Systems with Corules for Fair Subtyping and Liveness Properties of Binary Session Types. In *Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP'21) (LIPIcs, Vol. 198)*, Nikhil Bansal, Emanuela Merelli, and James Worrell (Eds.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 125:1–125:16. <https://doi.org/10.4230/LIPIcs.ICALP.2021.125>
- Byron Cook, Alexey Gotsman, Andreas Podelski, Andrey Rybalchenko, and Moshe Y. Vardi. 2007. Proving that programs eventually do something good. In *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007*, Martin Hofmann and Matthias Felleisen (Eds.). ACM, 265–276. <https://doi.org/10.1145/1190216.1190257>
- Bruno Courcelle. 1983. Fundamental Properties of Infinite Trees. *Theor. Comput. Sci.* 25 (1983), 95–169. [https://doi.org/10.1016/0304-3975\(83\)90059-2](https://doi.org/10.1016/0304-3975(83)90059-2)
- Francesco Dagnino. 2019. Coaxioms: flexible coinductive definitions by inference systems. *Log. Methods Comput. Sci.* 15, 1 (2019). [https://doi.org/10.23638/LMCS-15\(1:26\)2019](https://doi.org/10.23638/LMCS-15(1:26)2019)

- Ornela Dardha and Simon J. Gay. 2018. A New Linear Logic for Deadlock-Free Session-Typed Processes. In *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 10803)*, Christel Baier and Ugo Dal Lago (Eds.). Springer, 91–109. https://doi.org/10.1007/978-3-319-89366-2_5
- Ornela Dardha, Elena Giachino, and Davide Sangiorgi. 2017. Session types revisited. *Inf. Comput.* 256 (2017), 253–286. <https://doi.org/10.1016/j.ic.2017.06.002>
- Ornela Dardha and Jorge A. Pérez. 2015. Comparing Deadlock-Free Session Typed Processes. In *Proceedings of the Combined 22th International Workshop on Expressiveness in Concurrency and 12th Workshop on Structural Operational Semantics, EXPRESS/SOS 2015, Madrid, Spain, 31st August 2015 (EPTCS, Vol. 190)*, Silvia Crafa and Daniel Gebler (Eds.). 1–15. <https://doi.org/10.4204/EPTCS.190.1>
- Romain Demangeon, Daniel Hirschhoff, and Davide Sangiorgi. 2009. Mobile Processes and Termination. In *Semantics and Algebraic Specification, Essays Dedicated to Peter D. Mosses on the Occasion of His 60th Birthday (Lecture Notes in Computer Science, Vol. 5700)*, Jens Palsberg (Ed.). Springer, 250–273. https://doi.org/10.1007/978-3-642-04164-8_13
- Yuxin Deng and Davide Sangiorgi. 2006. Ensuring termination by typability. *Inf. Comput.* 204, 7 (2006), 1045–1082. <https://doi.org/10.1016/j.ic.2006.03.002>
- Amina Doumane. 2017. *On the infinitary proof theory of logics with fixed points. (Théorie de la démonstration infinitaire pour les logiques à points fixes)*. Ph. D. Dissertation. Paris Diderot University, France. <https://tel.archives-ouvertes.fr/tel-01676953>
- Sebastian Erdweg, Oliver Bracevac, Edlira Kuci, Matthias Krebs, and Mira Mezini. 2015. A co-contextual formulation of type rules and its application to incremental type checking. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2015, part of SPLASH 2015, Pittsburgh, PA, USA, October 25-30, 2015*, Jonathan Aldrich and Patrick Eugster (Eds.). ACM, 880–897. <https://doi.org/10.1145/2814270.2814277>
- Nissim Francez. 1986. *Fairness*. Springer. <https://doi.org/10.1007/978-1-4612-4886-6>
- Pierre Ganty, Rupak Majumdar, and Andrey Rybalchenko. 2009. Verifying liveness for asynchronous programs. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 102–113. <https://doi.org/10.1145/1480881.1480895>
- Simon J. Gay. 2016. Subtyping Supports Safe Session Substitution. In *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday (Lecture Notes in Computer Science, Vol. 9600)*, Sam Lindley, Conor McBride, Philip W. Trinder, and Donald Sannella (Eds.). Springer, 95–108. https://doi.org/10.1007/978-3-319-30936-1_5
- Simon J. Gay and Malcolm Hole. 2005. Subtyping for session types in the pi calculus. *Acta Informatica* 42, 2-3 (2005), 191–225. <https://doi.org/10.1007/s00236-005-0177-z>
- Simon J. Gay and Vasco Thudichum Vasconcelos. 2010. Linear type theory for asynchronous session types. *J. Funct. Program.* 20, 1 (2010), 19–50. <https://doi.org/10.1017/S0956796809990268>
- Orna Grumberg, Nissim Francez, and Shmuel Katz. 1984. Fair Termination of Communicating Processes. In *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing (Vancouver, British Columbia, Canada) (PODC '84)*, Association for Computing Machinery, New York, NY, USA, 254–265. <https://doi.org/10.1145/800222.806752>
- Kohei Honda. 1993. Types for Dyadic Interaction. In *CONCUR '93, 4th International Conference on Concurrency Theory, Hildesheim, Germany, August 23-26, 1993, Proceedings (Lecture Notes in Computer Science, Vol. 715)*, Eike Best (Ed.). Springer, 509–523. https://doi.org/10.1007/3-540-57208-2_35
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. 1998. Language Primitives and Type Discipline for Structured Communication-Based Programming. In *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 (Lecture Notes in Computer Science, Vol. 1381)*, Chris Hankin (Ed.). Springer, 122–138. <https://doi.org/10.1007/BFb0053567>
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty asynchronous session types. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, George C. Necula and Philip Wadler (Eds.). ACM, 273–284. <https://doi.org/10.1145/1328438.1328472>
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty Asynchronous Session Types. *J. ACM* 63, 1 (2016), 9:1–9:67. <https://doi.org/10.1145/2827695>
- Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniérou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. 2016. Foundations of Session Types and Behavioural Contracts. *ACM Comput. Surv.* 49, 1 (2016), 3:1–3:36. <https://doi.org/10.1145/2873052>
- Naoki Kobayashi. 2002. A Type System for Lock-Free Processes. *Inf. Comput.* 177, 2 (2002), 122–159. <https://doi.org/10.1006/inco.2002.3171>

- Naoki Kobayashi. 2006. A New Type System for Deadlock-Free Processes. In *CONCUR 2006 - Concurrency Theory, 17th International Conference, CONCUR 2006, Bonn, Germany, August 27-30, 2006, Proceedings (Lecture Notes in Computer Science, Vol. 4137)*, Christel Baier and Holger Hermanns (Eds.). Springer, 233–247. https://doi.org/10.1007/11817949_16
- Naoki Kobayashi and Cosimo Laneve. 2017. Deadlock analysis of unbounded process networks. *Inf. Comput.* 252 (2017), 48–70. <https://doi.org/10.1016/j.ic.2016.03.004>
- Naoki Kobayashi and Davide Sangiorgi. 2010. 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>
- M.Z. Kwiatkowska. 1989. Survey of fairness notions. *Information and Software Technology* 31, 7 (1989), 371–386. [https://doi.org/10.1016/0950-5849\(89\)90159-6](https://doi.org/10.1016/0950-5849(89)90159-6)
- Leslie Lamport. 2000. Fairness and hyperfairness. *Distributed Comput.* 13, 4 (2000), 239–245. <https://doi.org/10.1007/PL00008921>
- Sam Lindley and J. Garrett Morris. 2016. Talking bananas: structural recursion for session types. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 434–447. <https://doi.org/10.1145/2951913.2951921>
- Barbara Liskov and Jeannette M. Wing. 1994. A Behavioral Notion of Subtyping. *ACM Trans. Program. Lang. Syst.* 16, 6 (1994), 1811–1841. <https://doi.org/10.1145/197320.197383>
- V. Natarajan and Rance Cleaveland. 1995. Divergence and Fair Testing. In *Automata, Languages and Programming, 22nd International Colloquium, ICALP95, Szeged, Hungary, July 10-14, 1995, Proceedings (Lecture Notes in Computer Science, Vol. 944)*, Zoltán Fülöp and Ferenc Gécseg (Eds.). Springer, 648–659. https://doi.org/10.1007/3-540-60084-1_112
- Susan S. Owicki and Leslie Lamport. 1982. Proving Liveness Properties of Concurrent Programs. *ACM Trans. Program. Lang. Syst.* 4, 3 (1982), 455–495. <https://doi.org/10.1145/357172.357178>
- Luca Padovani. 2013. Fair Subtyping for Open Session Types. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 7966)*, Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg (Eds.). Springer, 373–384. https://doi.org/10.1007/978-3-642-39212-2_34
- Luca Padovani. 2014. Deadlock and lock freedom in the linear π -calculus. In *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*, Thomas A. Henzinger and Dale Miller (Eds.). ACM, 72:1–72:10. <https://doi.org/10.1145/2603088.2603116>
- Luca Padovani. 2016. Fair subtyping for multi-party session types. *Math. Struct. Comput. Sci.* 26, 3 (2016), 424–464. <https://doi.org/10.1017/S096012951400022X>
- Luca Padovani, Vasco Thudichum Vasconcelos, and Hugo Torres Vieira. 2014. Typing Liveness in Multiparty Communicating Systems. In *Coordination Models and Languages - 16th IFIP WG 6.1 International Conference, COORDINATION 2014, Held as Part of the 9th International Federated Conferences on Distributed Computing Techniques, DisCoTec 2014, Berlin, Germany, June 3-5, 2014, Proceedings (Lecture Notes in Computer Science, Vol. 8459)*, eva Kühn and Rosario Pugliese (Eds.). Springer, 147–162. https://doi.org/10.1007/978-3-662-43376-8_10
- Arend Rensink and Walter Vogler. 2007. Fair testing. *Inf. Comput.* 205, 2 (2007), 125–198. <https://doi.org/10.1016/j.ic.2006.06.002>
- Alceste Scalas and Nobuko Yoshida. 2019. Less is more: multiparty session types revisited. *Proc. ACM Program. Lang.* 3, POPL (2019), 30:1–30:29. <https://doi.org/10.1145/3290343>
- Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A Higher-Order Logic for Concurrent Termination-Preserving Refinement. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10201)*, Hongseok Yang (Ed.). Springer, 909–936. https://doi.org/10.1007/978-3-662-54434-1_34
- Rob van Glabbeek. 2019. Justness - A Completeness Criterion for Capturing Liveness Properties (Extended Abstract). In *Foundations of Software Science and Computation Structures - 22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11425)*, Mikolaj Bojanczyk and Alex Simpson (Eds.). Springer, 505–522. https://doi.org/10.1007/978-3-030-17127-8_29
- Rob van Glabbeek and Peter Höfner. 2019. Progress, Justness, and Fairness. *ACM Comput. Surv.* 52, 4 (2019), 69:1–69:38. <https://doi.org/10.1145/3329125>
- Rob van Glabbeek, Peter Höfner, and Ross Horne. 2021. Assuming Just Enough Fairness to make Session Types Complete for Lock-freedom. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*. IEEE, 1–13. <https://doi.org/10.1109/LICS52264.2021.9470531>
- Philip Wadler. 2014. Propositions as sessions. *J. Funct. Program.* 24, 2-3 (2014), 384–418. <https://doi.org/10.1017/S095679681400001X>

Nobuko Yoshida, Martin Berger, and Kohei Honda. 2004. Strong normalisation in the pi -calculus. *Inf. Comput.* 191, 2 (2004), 145–202. <https://doi.org/10.1016/j.ic.2003.08.004>