

A Universal Session Type for Untyped Asynchronous Communication

Stephanie Balzer¹

Carnegie Mellon University, USA

Frank Pfenning²

Carnegie Mellon University, USA

Bernardo Toninho³

NOVA LINCS, Universidade Nova de Lisboa, Portugal

Abstract

In the simply-typed λ -calculus we can recover the full range of expressiveness of the untyped λ -calculus solely by adding a single recursive type $\mathcal{U} = \mathcal{U} \rightarrow \mathcal{U}$. In contrast, in the session-typed π -calculus, recursion alone is insufficient to recover the untyped π -calculus, primarily due to linearity: each channel just has two unique endpoints. In this paper, we show that shared channels with a corresponding sharing semantics (based on the language SILL_S developed in prior work) are enough to embed the untyped asynchronous π -calculus via a universal shared session type \mathcal{U}_S . We show that our encoding of the asynchronous π -calculus satisfies operational correspondence and preserves observable actions (i.e., processes are weakly bisimilar to their encoding). Moreover, we clarify the expressiveness of SILL_S by developing an operationally correct encoding of SILL_S in the asynchronous π -calculus.

2012 ACM Subject Classification Theory of computation \rightarrow Process calculi, Theory of computation \rightarrow Linear logic

Keywords and phrases Session types, sharing, π -calculus, bisimulation

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2018.30

1 Introduction

Session types [20, 22, 23] prescribe the protocols of message exchange between processes that interact along channels. The recent discovery of a Curry-Howard isomorphism between *linear logic* and the *session-typed π -calculus* [8, 9, 42, 38] has given message-passing concurrency a firm logical foundation. Programming languages [40, 19] building on this isomorphism not only guarantee *session fidelity* (i.e., protocol compliance) but also a form of *global progress*, since the process graph forms a tree and is acyclic by construction.

While the linear logic session framework allows for persistent servers through the exponential modality (i.e., replicated sessions that may be used an arbitrary number of times), it enforces a strict separation between server instances by means of a copying semantics [8, 42]. For instance, interactions between a client and a server cannot affect future client-server interactions. Thus, this session discipline fundamentally excludes programming scenarios that require *sharing* of server resources such as shared databases or shared output devices. This observation triggered the realization that linear session-typed calculi lag behind the untyped

¹ NSF Grant No. CCF-1718267: “Enriching Session Types for Practical Concurrent Programming”

² NSF Grant No. CCF-1718267: “Enriching Session Types for Practical Concurrent Programming”

³ NOVA LINCS (Ref. UID/CEC/04516/2013)



© Stephanie Balzer, Frank Pfenning, and Bernardo Toninho;
licensed under Creative Commons License CC-BY

29th International Conference on Concurrency Theory (CONCUR 2018).
Editors: Sven Schewe and Lijun Zhang; Article No. 30; pp. 30:1–30:18



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

asynchronous π -calculus in expressiveness and the question of whether the full expressiveness of the untyped asynchronous π -calculus could be recovered in such a logical setting [42].

In this paper, we answer this question *positively*. In prior work we have introduced *manifest sharing* [3], a modal-typing discipline that orchestrates the coexistence of linear *and* shared channels while maintaining session fidelity, at the expense of generalized deadlock-freedom. In this work we show that manifest sharing recovers the expressiveness of the untyped asynchronous π -calculus. Given our language SILL_5 [3, 4] that supports manifest sharing, we provide an encoding of the untyped asynchronous π -calculus into SILL_5 , showing that our encoding satisfies *operational correspondence* and that π -calculus processes are *weakly bisimilar* to their SILL_5 encodings. To clarify the expressiveness of SILL_5 , we moreover develop an encoding in the other direction, embedding SILL_5 into the asynchronous (polyadic) π -calculus and satisfying operational correspondence.

Key to our encoding of the untyped asynchronous π -calculus into SILL_5 is the representation of a π -calculus channel as a recursive *shared* session type \mathcal{U}_s , reminiscent of the encoding of the untyped λ -calculus into the simply-typed λ -calculus via the type $\mathcal{U} = \mathcal{U} \rightarrow \mathcal{U}$. While the addition of a single recursive type is sufficient to recover the expressiveness of the untyped λ -calculus in the simply-typed λ -calculus, our result reveals that *both* shared and recursive session types are necessary to achieve the analogous result in the session-typed π -calculus.

The contributions of this paper are:

- A proof that our encoding of the untyped asynchronous π -calculus into SILL_5 is operationally sound and complete and preserves observable actions (i.e., processes are weakly bisimilar to their encoding);⁴
- A formulation of a weak bisimulation between a labelled transition system for the asynchronous π -calculus and a multiset rewriting system for *closed terms* of SILL_5 ;
- Evidence of the instrumental role shared channels take in the expressiveness of session-typed process calculi;
- An encoding of SILL_5 into the untyped asynchronous polyadic π -calculus, satisfying operational correspondence.

Paper Structure. Section 2 provides the necessary background on SILL_5 . Section 3 introduces the encoding of the untyped asynchronous π -calculus into SILL_5 and states and proves operational and observational correspondence (i.e., preservation of reductions and observable actions). Section 4 develops an encoding of SILL_5 into the untyped asynchronous polyadic π -calculus, satisfying operational correspondence. Section 5 summarizes related work, and Section 6 concludes the paper. Proofs are given in a companion technical report.

2 Manifest Sharing with Session Types

In this section, we provide an introduction to *manifest sharing* [3] and the programming language SILL_5 [3, 4], to the extent necessary for the development in this paper. *Session types* [20, 22, 23, 8, 40, 9, 42, 38] prescribe the protocols of message exchange between processes that interact along channels. For example, the recursive linear session type

$$\text{queue } A = \&\{\text{enq} : A \multimap \text{queue } A, \text{deq} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \otimes \text{queue } A\}\}$$

⁴ A preliminary version of our encoding of the untyped asynchronous π -calculus into SILL_5 has been published in [3] for illustration purposes, but without proof.

■ **Table 1** Overview of session types in SILL_5 together with their operational meaning.

Session type	Process term			Description
current	cont	current	cont	
$c_L : \oplus\{\overline{l} : A_L\}$	$c_L : A_{L_h}$	$c_L.l_h ; P$	P	provider sends label l_h along c_L
		$\text{case } c_L \text{ of } \overline{l} \Rightarrow Q$	Q_h	client receives label l_h along c_L
$c_L : \&\{\overline{l} : A_L\}$	$c_L : A_{L_h}$	$\text{case } c_L \text{ of } \overline{l} \Rightarrow P$	P_h	provider receives label l_h along c
		$c_L.l_h ; Q$	Q	client sends label l_h along c_L
$c_L : A_L \otimes B_L$	$c_L : B_L$	$\text{send } c_L d_L ; P$	P	provider sends channel $d_L : A_L$ along c_L
		$y_L \leftarrow \text{recv } c_L ; Q_{y_L}$	$[d_L/y_L] Q_{y_L}$	client receives channel $d_L : A_L$ along c_L
$c_L : A_L \multimap B_L$	$c_L : B_L$	$y_L \leftarrow \text{recv } c_L ; P_{y_L}$	$[d_L/y_L] P_{y_L}$	provider receives channel $d_L : A_L$ along c_L
		$\text{send } c_L d_L ; Q$	Q	client sends channel $d_L : A_L$ along c_L
$c_L : \Pi x:A_S.B_L$	$c_L : B_L$	$\text{send } c_L d_S ; P$	P	provider sends channel $d_S : A_S$ along c_L
		$y_S \leftarrow \text{recv } c_L ; Q_{y_S}$	$[d_S/y_S] Q_{y_S}$	client receives channel $d_S : A_S$ along c_L
$c_L : \exists x:A_S.B_L$	$c_L : B_L$	$y_S \leftarrow \text{recv } c_L ; P_{y_S}$	$[d_S/y_S] P_{y_S}$	provider receives channel $d_S : A_S$ along c_L
		$\text{send } c_L d_S ; Q$	Q	client sends channel $d_S : A_S$ along c_L
$c_L : \mathbf{1}$	-	$\text{close } c_L$	-	provider sends “end” along c_L
		$\text{wait } c_L ; Q$	Q	provider receives “end” along c_L
$c_L : \downarrow_L^S A_S$	$c_S : A_S$	$c_S \leftarrow \text{detach } c_L ; P_{x_S}$	$[c_S/x_S] P_{x_S}$	provider sends “detach c_S ” along c_L
		$x_S \leftarrow \text{release } c_L ; Q_{x_S}$	$[c_S/x_S] Q_{x_S}$	client receives “detach c_S ” along c_L
$c_S : \uparrow_L^S A_L$	$c_L : A_L$	$c_L \leftarrow \text{acquire } c_S ; Q_{x_L}$	$[c_L/x_L] Q_{x_L}$	client sends “acquire c_L ” along c_S
		$x_L \leftarrow \text{accept } c_S ; P_{x_L}$	$[c_L/x_L] P_{x_L}$	provider receives “acquire c_L ” along c_S

defines the protocol of how to interact with a provider of a queue data structure that contains elements of some variable type A . In a session-typed interpretation of intuitionistic linear logic, session types are expressed from the point of view of the providing process, with the channel along which the process provides the session behavior being defined by the session type. This choice avoids the explicit dualization of a session type present in the original presentations of session types [20, 22] and those based on classical linear logic [42]. We adopt an *equi-recursive* [11] interpretation for recursive session types, silently equating a recursive session type with its unfolding and requiring types to be *contractive* [16].

Table 1 provides an overview of SILL_5 's session types and their operational reading. For each type constructor, Table 1 lists the points of view of the *provider* and *client* of the given type, in the first and second lines, respectively. For each connective, its session type before the exchange (**Session type current**) and after the exchange (**Session type cont(inuation)**) is given. Likewise, the implementing process term is indicated before the exchange (**Process term current**) and after the exchange (**Process term continuation**). Table 1 shows that the process terms of a provider and a client for a connective come in matching pairs. Both participants' view of the session changes consistently.

For the linear session type *queue* A specified above, we have the following protocol: a process providing a service of type *queue* A gives a client the choice to either enqueue (**enq**) or dequeue (**deq**) an element of type A . Upon receipt of the label **enq**, the providing process expects to receive a channel of type A to be enqueued and recurs. Upon receipt of the label **deq**, the providing process either indicates that the queue is empty (**none**), in which case it terminates, or that there is a channel stored in the queue (**some**), in which case it dequeues this element, sends it to the client, and recurs.

Linearity restricts session type *queue* A to a single client. If we want the queue to be used in a classical consumer-producer scenario, where we have multiple producers and consumers

```

 $q' \leftarrow \text{acquire } q ;$ 
 $q'.\text{enq} ;$ 
 $\text{send } q' \ x ;$ 
 $q \leftarrow \text{release } q'$ 

```

■ **Figure 1** A client of a shared queue.

accessing the queue, we can use the following *shared session type* instead:

$$\mathbf{queue } A_s = \uparrow_L^s \&\{\text{enq} : \Pi x : A_s . \downarrow_L^s \mathbf{queue } A_s, \\ \text{deq} : \oplus\{\text{none} : \downarrow_L^s \mathbf{queue } A_s, \text{some} : \exists x : A_s . \downarrow_L^s \mathbf{queue } A_s\}\}$$

For ease of reading, we typeset shared session types and channels in **red** and **bold** font as opposed to linear session types and channels, which we typeset in black and regular font. Session type $\mathbf{queue } A_s$ now describes the session offered by a shared process. Since a shared process can have multiple clients that refer to the process by a *shared channel*, state-altering communication with a shared process must only happen once exclusive access to the process has been obtained. Otherwise, session fidelity would be endangered. To this end, SILL₅ imposes an *acquire-release* discipline on shared processes, where an acquire yields exclusive access to a shared process, if the process is available, and a release relinquishes exclusive access. As a result, processes can alternate between linear and shared, where a successful acquire of a shared process turns the process into a linear one, and conversely, a release of a linear process turns the process into a shared one.

A potential producer process can now interact with a process that implements session type $\mathbf{queue } A_s$ according to Figure 1, assuming that q is of type $\mathbf{queue } A_s$ and x is of type A_s . The statement $q' \leftarrow \text{acquire } q$, yields, if successful, the queue's linear channel q' along which the producer process can enqueue the element. The statement $q \leftarrow \text{release } q'$ releases the now linear \mathbf{queue} process providing along q' , giving turn to another producer or consumer process, and yields the queue's shared channel q . As indicated by Table 1, there exist the dual notions of an *accept* and *detach* for an acquire and release, respectively, denoting the matching statements by a provider.

A key contribution of manifest sharing is not only to support acquire-release as a programming primitive but also to make it manifest in the type system. Generalizing the idea of type *stratification* [35, 6, 36], session types are stratified into a linear and shared layer with two *adjoint modalities* going back and forth between them:

$$\begin{aligned} A_s &\triangleq \uparrow_L^s A_L \\ A_L, B_L &\triangleq A_L \otimes B_L \mid \mathbf{1} \mid \oplus\{\overline{l : A_L}\} \mid \exists x : A_s . B_L \mid A_L \multimap B_L \mid \Pi x : A_s . B_L \mid \&\{\overline{l : A_L}\} \mid \downarrow_L^s A_s \end{aligned}$$

The modal operator $\downarrow_L^s A_s$ shifting *down* from the shared to the linear layer is then interpreted as a *release* (and, dually, *detach*) and the operator $\uparrow_L^s A_L$ shifting *up* from the linear to the shared layer as an *acquire* (and, dually, *accept*). As a result, we obtain a type system where a session type dictates any form of synchronization, including the acquisition and release of a shared process.

Returning to the shared session type $\mathbf{queue } A_s$ defined above, we can see that any exchange of labels or channels with the queue is now guarded by a preceding acquire, and that the queue must be released before it recurs. The shared session type further deviates from its linear version in that it contains shared elements, as the entire queue is shared, and by recurring in the empty case of a dequeuing request, as there are now multiple clients.

We briefly discuss the typing and the dynamics of acquire-release. The typing and the dynamics of the residual linear connectives are standard. As is usual for an intuitionistic

$$\begin{array}{l}
\text{proc}(c_L, x_L \leftarrow \text{acquire } a_S; Q_{x_L}), \text{proc}(a_S, x_L \leftarrow \text{accept } a_S; P_{x_L}) \\
\longrightarrow \text{proc}(c_L, [a_L/x_L] Q_{x_L}), \text{proc}(a_L, [a_L/x_L] P_{x_L}), \text{unavail}(a_S) \\
\\
\text{proc}(c_L, x_S \leftarrow \text{release } a_L; Q_{x_S}), \text{proc}(a_L, x_S \leftarrow \text{detach } a_L; P_{x_S}), \text{unavail}(a_S) \\
\longrightarrow \text{proc}(c_L, [a_S/x_S] Q_{x_S}), \text{proc}(a_S, [a_S/x_S] P_{x_S})
\end{array}$$

■ **Figure 2** Multiset rewriting rules for acquire-release.

interpretation, each connective gives rise to a left and a right rule, denoting the use and provision, respectively, of a session of the given type:

$$\begin{array}{c}
\text{(T-}\uparrow\text{L)} \\
\frac{\Gamma, x_S : \uparrow\text{L}^S A_L; \Delta, x_L : A_L \vdash_{\Sigma} Q_{x_L} :: (z_L : C_L)}{\Gamma, x_S : \uparrow\text{L}^S A_L; \Delta \vdash_{\Sigma} x_L \leftarrow \text{acquire } x_S; Q_{x_L} :: (z_L : C_L)} \\
\text{(T-}\downarrow\text{L)} \\
\frac{\Gamma, x_S : A_S; \Delta \vdash_{\Sigma} Q_{x_S} :: (z_L : C_L)}{\Gamma; \Delta, x_L : \downarrow\text{L}^S A_S \vdash_{\Sigma} x_S \leftarrow \text{release } x_L; Q_{x_S} :: (z_L : C_L)}
\end{array}
\qquad
\begin{array}{c}
\text{(T-}\uparrow\text{R)} \\
\frac{\Gamma; \cdot \vdash_{\Sigma} P_{x_L} :: (x_L : A_L)}{\Gamma \vdash_{\Sigma} x_L \leftarrow \text{accept } x_S; P_{x_L} :: (x_S : \uparrow\text{L}^S A_L)} \\
\text{(T-}\downarrow\text{R)} \\
\frac{\Gamma \vdash_{\Sigma} P_{x_S} :: (x_S : A_S)}{\Gamma; \cdot \vdash_{\Sigma} x_S \leftarrow \text{detach } x_L; P_{x_S} :: (x_L : \downarrow\text{L}^S A_S)}
\end{array}$$

The typing judgments $\Gamma \vdash_{\Sigma} P :: (x_S : A_S)$ and $\Gamma; \Delta \vdash_{\Sigma} P :: (x_L : A_L)$ indicate that process P provides a service of session type A along channel x , given the typing of services provided by processes along the channels in typing contexts Γ (and Δ). Γ and Δ consist of hypotheses on the typing of shared and linear channels, respectively, where Γ is a structural and Δ a linear context. To allow for recursive process definitions, the typing judgment depends on a signature Σ that is populated with all process definitions prior to type-checking. The adjoint formulation forbids a shared process from depending on linear channels [3, 35]. Thus, when a shared session accepts an acquire and shifts to linear, it starts with an empty linear context.

Operationally, the dynamics of SILL_S is captured by *multiset rewriting rules* [10], which denote computation in terms of state transitions between configurations of processes. Multiset rewriting rules are local in that they only mention the parts of a configuration they rewrite. For acquire-release we have the rules of Figure 2.

Configuration states are defined by the predicates $\text{proc}(c_m, P)$ and $\text{unavail}(a_S)$. The former denotes a process with process term P providing along channel c_m , the latter a placeholder for a shared process providing along channel a_S that is currently not available. The above rule exploits the invariant that a process' providing channel a can come at one of two modes, a linear one, a_L , and a shared one, a_S . While the process is linear, it provides along a_L , while it is shared, along a_S . When a process shifts between modes, it switches between the two modes of its offering channel. This channel at the appropriate mode is substituted for the variables occurring in process terms.

3 Recovering the Untyped Asynchronous π -calculus in SILL_S

We now detail our encoding of the asynchronous π -calculus into SILL_S , show that it satisfies *operational correspondence* and that processes are *weakly bisimilar* to their SILL_S encodings.

3.1 Encoding the Untyped Asynchronous π -calculus in SILL_S

The essence of linear session-typed process calculi – treating channels as stateful resources – is fundamental in facilitating reasoning about session-typed programs and to guarantee strong properties, such as session fidelity and possibly deadlock-freedom. However, where channels in linear session-typed process calculi connect exactly one *sending* process with one *receiving* process, in the untyped π -calculus they may connect *multiple* sending and

$\text{empty} : \{\mathcal{U}_s\}$ $\mathbf{c} \leftarrow \text{empty} =$ $\mathbf{c}' \leftarrow \text{accept } \mathbf{c};$ $\text{case } \mathbf{c}' \text{ of}$ $ \text{ins} \rightarrow \mathbf{x} \leftarrow \text{rcv } \mathbf{c}';$ $\quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}';$ $\quad \mathbf{e} \leftarrow \text{empty};$ $\quad \mathbf{c} \leftarrow \text{elem} \leftarrow \mathbf{x}, \mathbf{e}$ $ \text{del} \rightarrow \mathbf{c}'.\text{none};$ $\quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}';$ $\quad \mathbf{c} \leftarrow \text{empty}$	$\text{nd_pick} : \{\exists x:\mathcal{U}_s. \downarrow_1^s \mathcal{U}_s \leftarrow \mathcal{U}_s, \mathcal{U}_s\}$ $\mathbf{c} \leftarrow \text{nd_pick} \leftarrow \mathbf{x}, \mathbf{d} =$ $\text{ndc} \leftarrow \text{nd_choice};$ $\text{case } \text{ndc} \text{ of}$ $ \text{yes} \rightarrow \text{send } \mathbf{c}' \mathbf{x};$ $\quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}';$ $\quad \text{wait } \text{ndc};$ $\quad \text{fwd } \mathbf{c} \mathbf{d}$ $ \text{no} \rightarrow \mathbf{d}' \leftarrow \text{acquire } \mathbf{d};$ $\quad \mathbf{d}'.\text{del};$ $\text{case } \mathbf{d}' \text{ of}$ $ \text{none} \rightarrow \mathbf{d} \leftarrow \text{release } \mathbf{d}';$ $\quad \text{send } \mathbf{c}' \mathbf{x};$ $\quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}';$ $\quad \text{wait } \text{ndc};$ $\quad \text{fwd } \mathbf{c} \mathbf{d}$ $ \text{some} \rightarrow \mathbf{y} \leftarrow \text{rcv } \mathbf{d}';$ $\quad \mathbf{d} \leftarrow \text{release } \mathbf{d}';$ $\quad \text{send } \mathbf{c}' \mathbf{y};$ $\quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}';$ $\quad \text{wait } \text{ndc};$ $\quad \mathbf{c} \leftarrow \text{elem} \leftarrow \mathbf{x}, \mathbf{d}$	$\text{nd_choice} : \{\oplus\{\text{yes} : \mathbf{1}, \text{no} : \mathbf{1}\}\}$ $\mathbf{d} \leftarrow \text{nd_choice} =$ $\mathbf{c} \leftarrow \text{coin_head};$ $\mathbf{f} \leftarrow \text{coin_flipper} \leftarrow \mathbf{c};$ $\mathbf{c}' \leftarrow \text{acquire } \mathbf{c};$ $\text{case } \mathbf{c}' \text{ of}$ $ \text{head} \rightarrow \mathbf{c} \leftarrow \text{release } \mathbf{c}';$ $\quad \mathbf{d}.\text{yes};$ $\quad \text{wait } \mathbf{f};$ $\quad \text{close } \mathbf{d}$ $ \text{tail} \rightarrow \mathbf{c} \leftarrow \text{release } \mathbf{c}';$ $\quad \mathbf{d}.\text{no};$ $\quad \text{wait } \mathbf{f};$ $\quad \text{close } \mathbf{d}$
$\text{elem} : \{\mathcal{U}_s \leftarrow \mathcal{U}_s, \mathcal{U}_s\}$ $\mathbf{c} \leftarrow \text{elem} \leftarrow \mathbf{x}, \mathbf{d} =$ $\mathbf{c}' \leftarrow \text{accept } \mathbf{c};$ $\text{case } \mathbf{c}' \text{ of}$ $ \text{ins} \rightarrow \mathbf{y} \leftarrow \text{rcv } \mathbf{c}';$ $\quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}';$ $\quad \mathbf{e} \leftarrow \text{elem} \leftarrow \mathbf{x}, \mathbf{d};$ $\quad \mathbf{c} \leftarrow \text{elem} \leftarrow \mathbf{y}, \mathbf{e}$ $ \text{del} \rightarrow \mathbf{c}'.\text{some};$ $\quad \mathbf{c}' \leftarrow \text{nd_pick} \leftarrow \mathbf{x}, \mathbf{d}$		

■ **Figure 3** Processes *empty* and *elem* implementing a π -calculus channel with auxiliary processes. See Figure 4 for processes *coin_head*, *coin_tail*, and *coin_flipper* and session type *coin*.

receiving processes, giving rise to *non-determinism*. For example, the π -calculus process $c(x).P \mid \bar{c}(a) \mid c(y).Q$, made up of three parallel components, where the first and third seek to input along channel c and the second outputs the name a along c , may reduce to either $[a/x]P \mid c(y).Q$ or $c(x).P \mid [a/y]Q$.

In purely linear session-typed process calculi, on the other hand, message exchange is completely *deterministic*, even in the presence of replicated or persistent sessions (this argument is made precise through a typed contextual equivalence for intuitionistic linear logic sessions in [34]). The addition of sharing to session-typed calculi – and with it non-determinism – suggests that it should now be possible to faithfully encode the untyped π -calculus. In previous work we have postulated this conjecture by providing an encoding of the untyped asynchronous π -calculus into SILL_5 [3], without any further proof. We now refine the encoding and prove it operationally and behaviorally correct.

The basic idea of our encoding is to represent a π -calculus *process* by a *linear* SILL_5 process and a π -calculus *channel* by a *shared* SILL_5 process. Reminiscent of the encoding of the untyped λ -calculus into the typed λ -calculus, we type π -calculus channels with a *universal recursive shared* session type \mathcal{U}_s :

$$\mathcal{U}_s = \uparrow_1^s \&\{\text{ins} : \Pi x:\mathcal{U}_s. \downarrow_1^s \mathcal{U}_s, \text{del} : \oplus\{\text{none} : \downarrow_1^s \mathcal{U}_s, \text{some} : \exists x:\mathcal{U}_s. \downarrow_1^s \mathcal{U}_s\}\}$$

Similar to the type *queue* A_s of Section 2, the type \mathcal{U}_s represents a buffer that stores elements, but with the elements being of type \mathcal{U}_s themselves and *without* maintaining any order. Figure 3 shows the processes *empty* and *elem* that implement session type \mathcal{U}_s . In SILL_5 , we declare the type of a defined process X with $X : \{A \leftarrow A_1, \dots, A_n\}$, indicating that the process provides a service of type A , using channels of type A_1, \dots, A_n . The definition of the process is then given by $x \leftarrow X \leftarrow y_1, \dots, y_n = P$, where P is the body of the process with occurrences of channels $y_1 : A_1, \dots, y_n : A_n$. A new process X providing along channel x is spawned with an expression of the form $x \leftarrow X \leftarrow y_1, \dots, y_n ; Q_x$, where Q_x is the continuation binding x . We refer to Table 1 for the meaning of the process terms.

$$\begin{array}{l}
\mathbf{coin} = \uparrow_l^s \oplus \{ \text{head} : \downarrow_l^s \mathbf{coin}, \\
\text{tail} : \downarrow_l^s \mathbf{coin} \} \quad \mathit{coin_head} : \{ \mathbf{coin} \} \quad \mathit{coin_tail} : \{ \mathbf{coin} \} \quad \mathit{coin_flipper} : \{ \mathbf{1} \leftarrow \mathbf{coin} \} \\
\mathbf{c} \leftarrow \mathit{coin_head} = \quad \mathbf{c} \leftarrow \mathit{coin_tail} = \quad d \leftarrow \mathit{coin_flipper} \leftarrow \mathbf{c} = \\
\mathbf{c}' \leftarrow \text{accept } \mathbf{c}; \quad \mathbf{c}' \leftarrow \text{accept } \mathbf{c}; \quad \mathbf{c}' \leftarrow \text{acquire } \mathbf{c}; \\
\mathbf{c}'.\text{head}; \quad \mathbf{c}'.\text{tail}; \quad \text{case } \mathbf{c}' \text{ of} \\
\mathbf{c} \leftarrow \text{detach } \mathbf{c}'; \quad \mathbf{c} \leftarrow \text{detach } \mathbf{c}'; \quad | \text{head} \rightarrow \mathbf{c} \leftarrow \text{release } \mathbf{c}'; \\
\mathbf{c} \leftarrow \mathit{coin_tail} \quad \mathbf{c} \leftarrow \mathit{coin_head} \quad \text{close } d \\
\quad \quad \quad | \text{tail} \rightarrow \mathbf{c} \leftarrow \text{release } \mathbf{c}'; \\
\quad \quad \quad \text{close } d
\end{array}$$

■ **Figure 4** Processes $\mathit{coin_head}$, $\mathit{coin_tail}$, and $\mathit{coin_flipper}$ and session type coin , upon which process $\mathit{nd_choice}$ in Figure 3 relies.

The buffer is implemented as a sequence of elem processes, ending in an empty process. The recursive process elem provides a buffer sequence along channel c and uses a channel $x : \mathcal{U}_s$ (the buffer element at the current position in the sequence) as well as a channel $d : \mathcal{U}_s$ (the next elem of the sequence). Process empty , on the other hand, provides an empty buffer sequence along channel c , without using any other channels. Both processes insert the received element at the head of the buffer sequence in the ins case, but handle the del case differently. Whereas process empty responds with label none , process elem responds with label some , followed by sending and deleting an *arbitrary* element from the buffer. Process elem achieves arbitrary deletion by recurring as process $\mathit{nd_pick}$. Process $\mathit{nd_pick}$, in turn, uses process $\mathit{nd_choice}$ to nondeterministically choose between sending and deleting the element at the current position in the sequence (case yes) or, possibly recursively, propagating the deletion request to the next element in the sequence (case no). While linear session-typed calculi are deterministic, *non-determinism* arises in \mathbf{SILL}_5 from the acquisition of shared channels, since it is unknown which client among all those competing to acquire a shared process will succeed. Process $\mathit{nd_choice}$ uses this fact and achieves non-determinism by reading a coin that it shares with process $\mathit{coin_flipper}$ (see Figure 4). Both processes then try to acquire the coin concurrently, which switches sides when read, with the result that the value read by $\mathit{nd_choice}$ depends on the order in which the coin is acquired.

Given the buffer abstraction, encoded π -calculus processes in \mathbf{SILL}_5 simply amount to “producers” and “consumers” of shared channels of type \mathcal{U}_s . Any such process can communicate along a π -calculus channel by acquiring the corresponding \mathbf{SILL}_5 channel of universal type. We are now ready to give the encoding of the untyped asynchronous monadic π -calculus [30, 37] into \mathbf{SILL}_5 . The syntax of the asynchronous π -calculus is [5]:

$$P \triangleq 0 \mid \bar{c}(a) \mid c(x).P \mid \nu c P \mid P_1 \mid P_2 \mid !P$$

0 denotes an inactive process. $\bar{c}(a)$ represents an asynchronous send of channel a along channel c . $c(x).P$ amounts to a guarded input, where the channel received along c is bound to x in the continuation P . $\nu c P$ introduces a new channel c that is bound in P . $P_1 \mid P_2$ denotes parallel composition of P_1 and P_2 , and $!P$ replication of P (i.e., an unbounded number of copies of P in parallel). We assume a standard reduction and labelled transition semantics, but where replication involves an explicit reduction (and τ transition) instead of expansion through structural congruence: $!P \longrightarrow P \mid !P$. Moreover, we enforce that structural congruence is only applied at the top-level of processes.

Our encoding, shown in Figure 5, yields for each π -calculus process P a corresponding linear process $\llbracket P \rrbracket$ in \mathbf{SILL}_5 , satisfying the typing judgment: $\Gamma_{\mathcal{F}}; \Gamma_{\mathcal{B}}; \Gamma_{\mathcal{I}}; \cdot \vdash_{\Sigma} \llbracket P \rrbracket :: (\cdot)$. We use an empty succedent to denote that the process does not provide any session. Since all communication is going to happen along π -calculus channels, i.e., the shared \mathbf{SILL}_5 processes of type \mathcal{U}_s , the linear \mathbf{SILL}_5 processes representing π -calculus processes merely become clients

$\llbracket 0 \rrbracket$	$= \cdot$	$snd : \{(\Pi x:\mathcal{U}_s. \mathbf{1}) \leftarrow \mathcal{U}_s\}$	$poll_rcv : \{(\exists x:\mathcal{U}_s. \mathbf{1}) \leftarrow \mathcal{U}_s\}$
$\llbracket \bar{c}(a) \rrbracket$	$= x \leftarrow snd \leftarrow \mathbf{c};$ send $x \ \mathbf{a};$ wait $x; \cdot$	$d \leftarrow snd \leftarrow \mathbf{c} =$ $\mathbf{x} \leftarrow rcv \ d;$ $\mathbf{c}' \leftarrow acquire \ \mathbf{c};$ $\mathbf{c}'.ins;$ send $\mathbf{c}' \ \mathbf{x};$ $\mathbf{c} \leftarrow release \ \mathbf{c}';$ close d	$d \leftarrow poll_rcv \leftarrow \mathbf{c} =$ $\mathbf{c}' \leftarrow acquire \ \mathbf{c};$ $\mathbf{c}'.del;$ case \mathbf{c}' of none $\rightarrow \mathbf{c} \leftarrow release \ \mathbf{c}';$ $d \leftarrow poll_rcv \leftarrow \mathbf{c}$ some $\rightarrow \mathbf{x} \leftarrow rcv \ \mathbf{c}';$ $\mathbf{c} \leftarrow release \ \mathbf{c}';$ send $d \ \mathbf{x};$ close d
$\llbracket c(x).P \rrbracket$	$= y \leftarrow poll_rcv \leftarrow \mathbf{c};$ $\mathbf{z} \leftarrow rcv \ y;$ wait $y;$ $\llbracket \mathbf{z}/x \rrbracket \llbracket P \rrbracket$	$\mathbf{c}'.ins;$ send $\mathbf{c}' \ \mathbf{x};$ $\mathbf{c} \leftarrow release \ \mathbf{c}';$ close d	$\mathbf{c}'.del;$ case \mathbf{c}' of none $\rightarrow \mathbf{c} \leftarrow release \ \mathbf{c}';$ $d \leftarrow poll_rcv \leftarrow \mathbf{c}$ some $\rightarrow \mathbf{x} \leftarrow rcv \ \mathbf{c}';$ $\mathbf{c} \leftarrow release \ \mathbf{c}';$ send $d \ \mathbf{x};$ close d
$\llbracket \nu x P \rrbracket$	$= \mathbf{y} \leftarrow empty;$ $\llbracket \mathbf{y}/x \rrbracket \llbracket P \rrbracket$		$\mathbf{c} \leftarrow release \ \mathbf{c}';$ send $d \ \mathbf{x};$ close d
$\llbracket P_1 \mid P_2 \rrbracket$	$= _ \leftarrow \llbracket P_1 \rrbracket;$ $\llbracket P_2 \rrbracket$		
$\llbracket !P \rrbracket$	$= Rec_{c,P}$ where		
$Rec_{c,P}$	$= _ \leftarrow \llbracket P \rrbracket;$ $Rec_{c,P}$		

■ **Figure 5** Translation of untyped asynchronous π -calculus processes into SILL₅ and auxiliary processes snd and $poll_rcv$ ($empty : \{\mathcal{U}_s\}$ is defined in Figure 3).

of those processes, without providing any behavior outright. In our earlier encoding [3], we have translated π -calculus processes into linear SILL₅ processes of type **1**, since the notion of a non-providing linear process is not present in SILL₅. Our current encoding avoids the spurious exchange of wait messages required by type **1** and constitutes a return to the original interpretation of linear logic [8], where processes terminate silently. In the above typing judgment, we moreover subdivide the context Γ into three parts, to keep track of the *free* ($\Gamma_{\mathcal{F}}$) and *bound* ($\Gamma_{\mathcal{B}}$) π -calculus channels as well as of channels that are only used *internally* to the encoding ($\Gamma_{\mathcal{I}}$). When an encoded process reduces, new linear channels may be generated, for example, the providing channel of process nd_choice , which are all internal to the encoding.

The inactive process 0 is encoded as the empty SILL₅ process. The encoding of an output $\llbracket \bar{c}(a) \rrbracket$ is implemented by spawning a new linear SILL₅ process snd of type $\Pi x:\mathcal{U}_s. \mathbf{1}$ with access to the buffer implementing channel c . The encoding then sends the channel a to the spawned process snd , waiting for snd to acquire the buffer c , insert a , and terminate. The encoding of an input $\llbracket c(x).P \rrbracket$ is implemented by spawning a new linear SILL₅ process $poll_rcv$ of type $\exists x:\mathcal{U}_s. \mathbf{1}$ with access to the buffer implementing channel c . The encoding then waits for the spawned process $poll_rcv$ to send back a channel and terminate, after which it continues at P , substituting the received channel for x . Process $poll_rcv$ repeatedly checks, in a potentially infinite loop, if the buffer c contains an element. If so, it deletes it from the buffer, passes it on, and terminates. New name creation ($\llbracket \nu x P \rrbracket$) simply spawns a new buffer, offering on some fresh name x . Parallel ($\llbracket P_1 \mid P_2 \rrbracket$) composition is embodied by a spawning of the processes P_1 in parallel with the executing process P_2 . Finally, replication ($\llbracket !P \rrbracket$) is implemented by a loop that spawns copies of the replicated process.

To make our encoding more tangible, we derive the initial SILL₅ configuration obtained from translating the process $\llbracket \bar{c}(a) \mid c(x).0 \rrbracket$ according to the rules in Figure 5:

$$\begin{aligned}
 a_s, c_s; \cdot; \cdot; \cdot \vdash_{\Sigma} & \text{proc}(_, y_L \leftarrow poll_rcv \leftarrow c_s; z_s \leftarrow rcv \ y_L; \text{wait } y_L; \cdot), \\
 & \text{proc}(_, y_L \leftarrow snd \leftarrow c_s; \text{send } y_L \ a_s; \text{wait } y_L; \cdot), \\
 & \text{buf}(a_s \mid y_L \leftarrow \text{accept } a_s; P_{y_L}), \text{buf}(c_s \mid y_L \leftarrow \text{accept } c_s; P_{y_L})
 \end{aligned}$$

To the left, we list the contents of the contexts $\Gamma_{\mathcal{F}}$; $\Gamma_{\mathcal{B}}$; $\Gamma_{\mathcal{I}}$; Δ , to the right the process configuration. For readability we use the short-form $\text{buf}(a \mid P_a)$ to represent a sequence of *empty*-terminated *elem* processes denoting an entire buffer, with P_a standing for the next

statements to be executed. The above configuration will reduce, according to the semantics of SILL_5 , until it halts in a state that consists of buffers representing the π -calculus channels, *coin_head* processes for any nondeterministic choices made, and *unavail* predicates for any shared channels that are not available. On the other hand, any linearly spawned processes that are internal to the encoding and not part of a buffer will have terminated.

Asynchrony of π -calculus outputs is achieved in our encoding by the introduction of the buffers, which temporarily store outputs until there is a process that is willing to receive. As a matter of fact, our buffers can be thought of manifestations of the “ether” to which asynchronous outputs are sent in the untyped asynchronous π -calculus! Our encoding is thus reminiscent of the encoding of the untyped asynchronous π -calculus into an untyped synchronous π -calculus with bags [5]. In fact, unlike the π -calculus where synchronous and asynchronous calculi have different expressive power [33], in the session-typed setting we can easily and selectively implement one in the other either by using double shifts to force acknowledgments [35] or by spawning single-message processes to achieve asynchrony [3]. The only significant point in SILL_5 is that acquire/accept interactions must be a synchronization point. As we discuss in Section 3.3, crucial to the correctness of our encoding is also the removal of buffer elements non-deterministically. This guarantees that at no point in a reduction is the order between outputs determined. The use of nondeterministic deletion is another improvement over our earlier encoding [3], which uses non-deterministic insertion.

An interested reader may wonder whether asynchronous messages could not be encoded directly as processes, rather than storing them temporarily in a buffer until their receipt. After all, this is exactly what the syntax of the asynchronous π -calculus enforces! Non-determinism would then be achieved by the operational dynamics of the multiset rewriting rules, eliminating the need for the explicit encoding of non-deterministic buffers. Since every π -calculus channel c is mapped to a shared SILL_5 channel c_s , this hypothetical encoding would require the ability to have multiple processes offering along the same shared channel (either the sender or the receiver sides of the communication). This is not allowed by the typing discipline, which crucially enforces that every process offers along a unique channel. Thus, an explicit representation of buffers is key, which then requires the encoding of non-deterministic bags to mimic the semantics of asynchrony in a precise way.

3.2 Operational Correspondence

We now develop an operational correspondence result for our encoding of the untyped asynchronous π -calculus. Operational correspondence results are standard *desiderata* for encodings of process calculi [18], showing that the computational features of the source language are preserved by the encoding in a precise sense. Following the terminology of [18], we aim to establish operational *completeness* (i.e., that π -calculus reductions are mimicked by the encoding) and *soundness* (i.e., that computations of encoded processes can be mapped back to those of the source terms) of our encoding.

As is the case in most encodings, some of the computation steps in the image of our encoding are purely administrative artifacts, and thus may not have a counterpart in the source. Specifically, the encoding of π -calculus channels as buffers introduces quite a few such “spurious” steps. Rather than relating source and image of the encoding at every step [5, 18], we introduce the notion of an *administrative* transition, and then state operational correspondence modulo such administrative transitions.

Given the nature of the asynchronous π -calculus, in which outputs are sent into the “ether” and synchronization only happens upon receipt, we deem the interactions leading to the insertion into a buffer as administrative and only the removal itself *relevant*. This treatment

is consistent with the existing literature. In the encoding of the untyped asynchronous π -calculus into an untyped synchronous π -calculus with bags [5], output prefixes are equated with one-element bags, and synchronization amounts to directly reading from these bags. We define relevant and administrative transitions in the image of our encoding as follows:

► **Definition 1** (Relevant and Administrative Transitions of Encoding). We say that a relevant transition, written \longrightarrow_r , is a standard transition between SILL configurations such that: $\Omega, \text{proc}(d_L, x_S \leftarrow \text{recv } c_L; Q_{x_S}) \longrightarrow \Omega', \text{proc}(d_L, [a_S/x_S] Q_{x_S})$, for some $\Omega, \Gamma_{\mathcal{F}}, \Gamma_{\mathcal{B}}, \Gamma_{\mathcal{I}}, a_S, c_S$, and d_L such that $a_S \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}, c_S \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}$, and $d_S \in \Gamma_{\mathcal{I}}$.

An administrative transition, written \longrightarrow_a , is a transition defined by the standard transition relation between SILL configurations, but excluding a relevant transition. We write \Longrightarrow_a for the reflexive transitive closure of \longrightarrow_a , and write \Longrightarrow_r for $\Longrightarrow_a \longrightarrow_r \Longrightarrow_a$.

Inspecting our encoding (Figure 3 and Figure 5), we can see that a relevant transition amounts to the receive action in the some branch in process poll_rcv , which synchronizes with the buffer to receive a channel. The parameters of the above definition uniquely identify this synchronization point: process poll_rcv is a linear process providing along a linear channel d_L that is internal to the encoding ($d_L \in \Gamma_{\mathcal{I}}$), and both the received channel a_S and the offering channel c_S of the buffer are either free or bound names of the original π -calculus process ($a_S \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}$ and $c_S \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}$).

Equipped with these two notions of transition, we can establish operational soundness and completeness. Their statements rely on the definition $\llbracket fn(P) \rrbracket$, which stand for a *configuration* of *empty* buffer processes of the form $\text{buf}(c_{s_i} \mid y_L \leftarrow \text{accept } c_{s_i}; Q_{y_L}), \dots, \text{buf}(c_{s_n} \mid y_L \leftarrow \text{accept } c_{s_n}; Q_{y_L})$, where $fn(P) = \{c_1, \dots, c_n\}$ denotes the set of free names in P . The definition allows us to compose an encoded π -calculus process with the appropriate buffer representations for all its free channel names.

► **Theorem 2** (Operational Correspondence).

Completeness. For all $P \longrightarrow P'$, there exists Ω_1, Ω_2 such that $\llbracket fn(P) \rrbracket, \text{proc}(_, \llbracket P \rrbracket) \Longrightarrow_r \Omega_1, \Omega_2$ or $\llbracket fn(P) \rrbracket, \text{proc}(_, \llbracket P \rrbracket) \Longrightarrow_a \Omega_1, \Omega_2$, with $\llbracket fn(P') \rrbracket, \text{proc}(_, \llbracket P' \rrbracket) \Longrightarrow_a \Omega_2$.

Soundness. For all P and $\llbracket fn(P) \rrbracket, \text{proc}(_, \llbracket P \rrbracket) \Longrightarrow_r \Omega$, there exists a P', Ω_1, Ω_2 such that $P \longrightarrow P'$ and $\Omega = \Omega_1, \Omega_2$ and $\llbracket fn(P') \rrbracket, \text{proc}(_, \llbracket P' \rrbracket) \Longrightarrow_a \Omega_2$.

For operational completeness, we identify each individual π -calculus reduction with either one relevant transition (possibly preceded or followed by several administrative transitions), or, for the π -calculus reduction corresponding to forking a parallel replica (i.e., $!P \longrightarrow P \mid !P$), with one administrative transition. For operational soundness, we match relevant transitions of encoded processes with one process reduction. In both settings we identify the artifacts of the encoding (coin processes and *unavail* channels) through the configuration Ω_1 .

We note that the encodings of continuations eventually “catch up” (via administrative transitions) with the configuration that results from the relevant transition, instead of having a more immediate identification through the encoding. This treatment is due to the distinction between processes (static entities) and configurations (runtime entities) in SILL_S, a distinction not present in the π -calculus, where processes *are* the runtime entities. For instance, parallel composition in SILL_S is achieved via an explicit spawning construct, whose semantics is to administratively *transition* to a configuration with the spawned process executing in parallel.

3.3 Observational Correspondence

In the previous section we have established that our encoding preserves reductions in the π -calculus in a strong sense, by identifying precisely the transitions in the operational semantics of SILL_5 that correspond to reductions in the π -calculus processes in a way that is consistent with standard results on the nature of asynchrony of the untyped asynchronous π -calculus.

We now go further and relate *observable actions* (i.e., labelled transitions) in the π -calculus with their corresponding observables in SILL_5 configuration rewrites. The key challenge here is to identify what those observables in SILL_5 are because of the significant differences between the semantic frameworks of the π -calculus and SILL_5 . Whereas the π -calculus adopts an *open-world* view of observable actions with an unspecified environment (the “ether”), SILL_5 adopts a *closed-world* view of a configuration of processes that are composed to form a complete program that can be run.

To clarify, consider the π -calculus process $\bar{c}\langle a \rangle \mid c(x).P$, where both c and a are free names. This process can interact with the environment through its free names by taking any of the following three observable actions: the output along c , the input along c , or the τ -action, corresponding to the synchronization between these dual actions. Now consider the SILL_5 encoding of $\llbracket \bar{c}\langle a \rangle \mid c(x).P \rrbracket$. It results in a complete configuration consisting of the encoding of the process together with an *explicit* encoding of the free names c and a in terms of the buffers offering along c and a . Given this setup, any potential action on the π -calculus side will result in a series of actual computational steps on the SILL_5 side, affecting the buffers as prescribed by the protocol of type \mathcal{U}_5 . In such a closed-world setting, trying to exactly mimic potential actions seems unnatural, if not impossible.

However, it is still the case that we want to relate π -calculus behavior with SILL_5 behavior in a precise sense. To reconcile the open-world view of a labelled transition semantics with the closed-world view of computational steps, we note that the encoding already accounts for this issue by essentially *implementing* “the environment” through the channel encodings that must be composed with the processes at the top-level. Thus, what we deem to be *observable* when we consider a configuration made up of encoded π -calculus processes and corresponding channel encodings are precisely the inputs and outputs to and from buffers. Conversely, any steps in a SILL_5 configuration that do not involve any inputs or outputs to and from buffers, we deem to be *unobservable*.

► **Definition 3** (Unobservable Transitions of Configuration). Given a configuration Ω we say that there is an unobservable transition from Ω to Ω' , written $\Omega \longrightarrow_{un} \Omega'$, iff $\Omega \longrightarrow \Omega'$ where the transition does *not* involve any of the two reductions below:

$$\begin{aligned} \Omega_0, \text{proc}(d_L, x_s \leftarrow \text{recv } c_L ; P_{x_s}) &\longrightarrow \Omega'_0, \text{proc}(d_L, [a_s/x_s] P_{x_s}) \\ \Omega_0, \text{proc}(d_L, \text{send } c_L e_s ; P) &\longrightarrow \Omega'_0, \text{proc}(d_L, P) \end{aligned}$$

for some $\Omega_0, \Omega'_0, \Gamma_{\mathcal{F}}, \Gamma_{\mathcal{B}}, \Gamma_{\mathcal{I}}, a_s, e_s, d_L$ and e_s such that $a_s, e_s, c_s \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}$, and $d_s \in \Gamma_{\mathcal{I}}$. We write $\Omega \Longrightarrow_{un} \Omega'$ to stand for the reflexive transitive closure of \longrightarrow_{un} .

► **Definition 4** (Observable Transitions of Configuration). Given a configuration Ω we define a notion of an observable transition $\Omega \xrightarrow{\alpha} \Omega'$, stating that configuration Ω performs action α and transitions to configuration Ω' , with $\alpha ::= \bar{c}\langle a \rangle \mid c(a) \mid (\nu a)\bar{c}\langle a \rangle \mid \tau$ as follows:

- $\Omega \xrightarrow{\bar{c}\langle a \rangle} \Omega'$ if $c, a \in \Gamma_{\mathcal{F}}$, $\Omega = \Omega_1, \text{proc}(d_L, \text{send } c a ; P), \Omega_2$, for some Ω_1, P, Ω_2 and $d_s \in \Gamma_{\mathcal{I}}$ and $\Omega \longrightarrow \Omega'$ with $\Omega' = \Omega'_1, \text{proc}(d_L, P), \Omega'_2$, for some Ω'_1, Ω'_2 .
- $\Omega \xrightarrow{(\nu a)\bar{c}\langle a \rangle} \Omega'$ if $c \in \Gamma_{\mathcal{F}}$, $a \in \Gamma_{\mathcal{B}}$, $\Omega = \Omega_1, \text{proc}(d_L, \text{send } c a ; P), \Omega_2$, for some Ω_1, P, Ω_2 and $d_s \in \Gamma_{\mathcal{I}}$ and $\Omega \longrightarrow \Omega'$ with $\Omega' = \Omega'_1, \text{proc}(d_L, P), \Omega'_2$, for some Ω'_1, Ω'_2 .

- $\Omega \xrightarrow{c(a)} \Omega'$ if $c \in \Gamma_{\mathcal{F}}$, $a \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}$, $\Omega = \Omega_1, \text{proc}(d_L, x \leftarrow \text{recv } c; P_x), \Omega_2$, for some Ω_1, P, Ω_2 and d , with $d_s \in \Gamma_{\mathcal{T}}$ and $\Omega \longrightarrow \Omega'$ with $\Omega' = \Omega'_1, \text{proc}(d_L, [a/x] P_x), \Omega'_2$, for some Ω'_1, Ω'_2 .
 - $\Omega \xrightarrow{\tau} \Omega'$ if all of the following:
 1. $c \in \Gamma_{\mathcal{B}}$, $a \in \Gamma_{\mathcal{F}} \cup \Gamma_{\mathcal{B}}$, $\Omega = \Omega_1, \text{proc}(d_L, \text{send } c \ a; P), \Omega_2$, for some Ω_1, P, Ω_2 and $d_s \in \Gamma_{\mathcal{T}}$ and $\Omega \longrightarrow \Omega''$ with $\Omega'' = \Omega'_1, \text{proc}(d_L, P), \Omega'_2$, for some $\Omega'_1, \Omega'_2, \Omega''$;
 2. $\Omega'' \Longrightarrow_{un} \Omega'''$ with $\Omega''' = \Omega''_1, \text{proc}(d_L, x \leftarrow \text{recv } c; Q_x), \Omega''_2$, for some $\Omega''_1, Q, \Omega''_2$ and d , with $d_s \in \Gamma_{\mathcal{T}}$ and $\Omega''' \longrightarrow \Omega'$ with $\Omega' = \Omega'_1, \text{proc}(d_L, [a/x] Q_x), \Omega'_2$, for some Ω'_1, Ω'_2 .
- We write $\Omega \xrightarrow{\alpha} \Omega'$ for $\Omega \Longrightarrow_{un} \Omega'' \xrightarrow{\alpha} \Omega''' \Longrightarrow_{un} \Omega'$.

The several observable transitions mirror the π -calculus labelled transitions, but where the role of the environment is replaced with the respective channel implementations. The first three cases define, respectively, output of a free name, output of a bound name, and input of a name (using the techniques of Definition 1 to track names). To account for synchronizations (τ -actions) in the π -calculus, we model the three steps that are required to perform a full communication in the encoding: an output action of a free or bound name to a buffer, followed by some sequence of unobservable transitions (needed to complete the several intermediate stages of the encoding), and an input action from the same buffer. With the right definition of observable in place, we define the natural notion of (weak) bisimulation between a π -calculus process and a SILL_S configuration.

► **Definition 5** (Weak Bisimulation). A relation \mathcal{R} between asynchronous π -calculus processes and SILL_S configurations is a weak bisimulation if and only if, whenever $PR\Omega$:

- If $P \xrightarrow{\alpha} P'$ and $\alpha \neq \tau$ then $\Omega \xrightarrow{\alpha} \Omega'$ and $P'\mathcal{R}\Omega'$
- If $P \xrightarrow{\tau} P'$ then $\Omega \Longrightarrow_{un} \Omega'$ or $\Omega \xrightarrow{\tau} \Omega'$ and $P'\mathcal{R}\Omega'$

plus the symmetric cases. We say that P is weakly bisimilar to Ω , written $P \approx \Omega$ iff there exists a weak bisimulation \mathcal{R} such that $PR\Omega$.

► **Theorem 6** (Observational Correspondence). *Let P be an asynchronous π -calculus process. We have that $P \approx \Omega, \text{proc}(_, \llbracket P \rrbracket)$, where Ω is a configuration made up of process encodings for the free names of P , with (non-empty) arbitrary contents.*

The expert reader may wonder how our use of a weak bisimulation captures asynchrony in the appropriate way, noting that a weak *asynchronous* bisimulation is necessary to accurately relate the asynchronous π -calculus and synchronous π -calculus with bags [5]. Would it then not be the case that we could use queues or stacks as buffers and replicate our bisimulation argument? Our argument holds *precisely* because of the non-deterministic (i.e., bag-like) nature of our buffer implementations. Otherwise, out-of-order message reception – a defining characteristic of asynchrony – would *not* be simulated correctly by our encoding. In this sense, our bisimulation is implicitly asynchronous by implementing the environment in terms of buffers that enforce non-deterministic removals.

4 Simulating Shared Session Types in the π -calculus

In this section, we close the loop and provide an encoding of SILL_S process terms into the asynchronous *polyadic* π -calculus. The extension to the polyadic π -calculus is necessary to send along with the actual channel a fresh continuation channel that must be used for the next exchange in the protocol. This continuation-passing-style encoding (similar to that of Dardha et al. [13]) ensures that messages are received in the order specified by the protocol.

The resulting encoding is summarised in Figure 6. To simplify our encoding, we use a type-directed expansion of forwarding corresponding to the standard identity expansion in the sequent calculus. The resulting programs no longer use forwarding as a primitive, but implement it by processes that forward messages from client to provider and vice versa. Observational correctness of this expansion has been shown for the linear fragment [7] and with recursive types [17]. The strong logical underpinnings lead us to conjecture that observational correctness extends to sharing as well.

The general pattern of the encoding is to translate a positive type [35] to an output and a negative type [35] to an input with matching bindings. In case of a linear output or input, a fresh continuation channel is provided in addition to the actual channel to be sent or received, respectively. This channel is then used in the process continuation (in parallel) in place of the original channel, guaranteeing that the session discipline is not disturbed by out-of-order messages. To encode the acquire-release discipline of SILL_5 , we must preserve the shared mode of a channel throughout the translation. To this end, we indicate a linear SILL_5 channel by a pair $\ulcorner x_L, x_S \urcorner$, where the left and right projections yield the linear mode x_L and shared mode x_S , respectively. A release then restores the session to the shared channel. To ensure a *blocking* semantics for an acquire, the encoding of an acquire and accept forces synchronization via the channel w . The encoding of choice makes use of a selection channel per choice, used to indicate the choice outcome and unlock the appropriate continuation. For simplicity, and without loss of generality, we limit the encoding to binary internal and external choice. Process definitions are encoded as top-level replicated processes:

For each $(x_L \leftarrow p \leftarrow \overline{y_L}, \overline{w_{S_j}} = P_{x_L, \overline{y_L}, \overline{w_{S_j}}}) \in \Sigma$:

$$!(p(\overline{y_L}, \overline{y_{S_i}}, \overline{w_{S_j}}, z) \cdot \nu x_L, x_S (\overline{z}(x_L, x_S) \mid \llbracket \ulcorner y_L, y_{S_i} \urcorner / \overline{y_L}, \ulcorner x_L, x_S \urcorner / x_L \rrbracket P_{x_L, \overline{y_L}, \overline{w_{S_j}}} \rrbracket))$$

For each $(x_S \leftarrow p \leftarrow \overline{y_{S_i}} = P_{x_S, \overline{y_{S_i}}}) \in \Sigma$: $!(p(\overline{y_{S_i}}, z) \cdot \nu x_S (\overline{z}(x_S) \mid \llbracket P_{x_S, \overline{y_{S_i}}} \rrbracket))$

The name of the definition is used as a channel that the encoding of the spawn construct uses to access new instances of the definition (generated via replication). The process receives the sessions that are needed to execute the definition and a channel z , used to send back the pair of (fresh) channels x_S and x_L used by the encoding of the definition body.

Operational Correspondence. To establish the operational correctness of our encoding, we consider an *asynchronous* semantics for SILL_5 . While operational completeness would not be affected by a synchronous semantics, soundness would require reasoning up-to observational equivalence. Since the expressiveness of SILL_5 has been shown to be orthogonal to the choice of synchrony or asynchrony, we opt for the latter for the sake of simplicity. The semantics spawns single-message outputting processes using a continuation-passing style to achieve type-safe asynchrony [3].

Recalling that in SILL_5 static entities are distinct from runtime entities, we lift the encoding to configurations, where the channels along which processes offer their session behavior are represented as bound names:

$$\llbracket \cdot \rrbracket = \mathbf{0} \quad \llbracket \text{proc}(c, P), \Omega \rrbracket = (\nu c_S, c_L)(\llbracket P \rrbracket \mid \llbracket \Omega \rrbracket) \quad \llbracket \text{unavail}(c_S), \Omega \rrbracket = \llbracket \Omega \rrbracket$$

We can now show that SILL_5 transitions are always matched by a synchronization in the π -calculus (and vice-versa) rather straightforwardly, given the direct nature of the encoding.

► **Theorem 7** (Operational Correspondence). *Let \longrightarrow^+ be the transitive closure of \longrightarrow :*

Completeness. *If P is a well-typed, forwarding-free SILL_5 process and $\text{proc}(a, P) \longrightarrow^+ \Omega$ then $\llbracket P \rrbracket \longrightarrow^+ \llbracket \Omega \rrbracket$.*

Soundness. *For all well-typed, forwarding-free SILL_5 configurations Ω such that $\llbracket \Omega \rrbracket \longrightarrow^+ Q$, there exists a configuration Ω' such that $\Omega \longrightarrow^+ \Omega'$ and $Q \Longrightarrow \llbracket \Omega' \rrbracket$.*

$$\begin{array}{ll}
\llbracket x_L \leftarrow p \leftarrow \overline{\Gamma y_{L_i}, y_{S_i}^\neg, w_{S_j}}; Q_{x_L} \rrbracket^{\text{SPAWN}_{LL}} & = \nu z (\overline{p} \langle \overline{y_{L_i}}, \overline{y_{S_i}}, \overline{w_{S_j}}, z \rangle \mid z(x_L, x_S). \llbracket \Gamma x_L, x_S^\neg / x_L \rrbracket Q_{x_L} \rrbracket) \\
\llbracket x_S \leftarrow p \leftarrow \overline{y_{S_i}}; Q_{x_S} \rrbracket^{\text{SPAWN}_{LL/SS}} & = \nu z (\overline{p} \langle \overline{y_{S_i}}, z \rangle \mid z(x_S). \llbracket Q_{x_S} \rrbracket) \\
\llbracket y_L \leftarrow \text{acquire } x_S; Q_{y_L} \rrbracket^{\uparrow_{LL}^S} & = \nu y_L, w (\overline{x_S} \langle y_L, x_S, w \rangle \mid w(\cdot). \llbracket \Gamma y_L, x_S^\neg / y_L \rrbracket Q_{y_L} \rrbracket) \\
\llbracket y_L \leftarrow \text{accept } x_S; P_{y_L} \rrbracket^{\downarrow_{LR}^S} & = x_S(y_L, y_S, w). (\overline{w} \langle \cdot \rangle \mid \llbracket \Gamma y_L, y_S^\neg / y_L \rrbracket \llbracket P_{y_L} \rrbracket) \\
\llbracket y_S \leftarrow \text{release } \Gamma x_L, x_S^\neg; Q_{y_S} \rrbracket^{\downarrow_{LL}^S} & = x_L(y_S). \llbracket Q_{y_S} \rrbracket \\
\llbracket x_S \leftarrow \text{detach } \Gamma x_L, x_S^\neg; P \rrbracket^{\downarrow_{LR}^S} & = \overline{x_L} \langle x_S \rangle \mid \llbracket P \rrbracket \\
\llbracket \text{wait } \Gamma x_L, x_S^\neg; Q \rrbracket^{\uparrow_{LL}^S} & = x_L(\cdot). \llbracket Q \rrbracket \\
\llbracket \text{close } \Gamma x_L, x_S^\neg \rrbracket^{\uparrow_{LR}^S} & = \overline{x_L} \langle \cdot \rangle \\
\llbracket y_L \leftarrow \text{recv } \Gamma x_L, x_S^\neg; P_{y_L} \rrbracket^{\otimes_{L/\circ R}} & = x_L(y_L, y_S, z_L, z_S). \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg, \Gamma y_L, y_S^\neg / y_L \rrbracket P_{y_L} \rrbracket \\
\llbracket \text{send } \Gamma x_L, x_S^\neg \Gamma y_L, y_S^\neg; P \rrbracket^{\otimes_{R/\circ L}} & = \nu z_L (\overline{x_L} \langle y_L, y_S, z_L, x_S \rangle \mid \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket P \rrbracket) \\
\llbracket y_S \leftarrow \text{recv } \Gamma x_L, x_S^\neg; P_{y_S} \rrbracket^{\exists_{LR}/\Pi R} & = x_L(y_S, z_L, z_S). \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket P_{y_S} \rrbracket \\
\llbracket \text{send } \Gamma x_L, x_S^\neg y_S; P \rrbracket^{\exists_{LR}/\Pi L} & = \nu z_L (\overline{x_L} \langle y_S, z_L, x_S \rangle \mid \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket P \rrbracket) \\
\llbracket \Gamma x_L, x_S^\neg. \text{case}(P, Q) \rrbracket^{\oplus_{L/\&R}} & = \nu y_{inl}, y_{inr} (\overline{x_L} \langle y_{inl}, y_{inr} \rangle \mid \\
& \quad y_{inl}(z_L, z_S). \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket P \rrbracket \mid \\
& \quad y_{inr}(z_L, z_S). \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket Q \rrbracket) \\
\llbracket \Gamma x_L, x_S^\neg. \text{inl}; P \rrbracket^{\oplus_{R1}/\&L1} & = \nu z_L (x_L(y_{inl}, y_{inr}). \overline{y_{inl}} \langle z_L, x_S \rangle \mid \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket P \rrbracket) \\
\llbracket \Gamma x_L, x_S^\neg. \text{inr}; Q \rrbracket^{\oplus_{R2}/\&L2} & = \nu z_L (x_L(y_{inl}, y_{inr}). \overline{y_{inr}} \langle z_L, x_S \rangle \mid \llbracket \llbracket \Gamma z_L, z_S^\neg / \Gamma x_L, x_S^\neg \rrbracket Q \rrbracket)
\end{array}$$

■ **Figure 6** Translation of SILL₅ process terms into the asynchronous, polyadic π -calculus.

5 Related Work

Encodings of Asynchrony. Encodability results are a standard benchmark for expressiveness of π -calculi [18]. For the asynchronous π -calculus [21], encodings into various formulations of synchronous π -calculi exist [5], as well as impossibility results [33] regarding the ability to adequately encode certain forms of choice in an asynchronous setting.

Our encoding of the asynchronous π -calculus is reminiscent of the encoding of the asynchronous π -calculus in a π -calculus with bags by Beauxis et al. [5], shown to be in tight correspondence via an asynchronous bisimilarity. Their framework considers buffers as primitives in the target calculus, whereas we encode the bag-like behavior of buffers explicitly as SILL₅ processes that adhere to a particularly typed protocol, making our encoding more primitive, but adding several administrative reductions to encoded processes due to the sharing discipline and the implementation of nondeterminism when reading from a buffer. This fact, combined with the restrictive (typed) usage of buffers in our setting allows us to reason using a weak bisimilarity rather than a more involved *asynchronous* bisimilarity. Beauxis et al. also consider an encoding of their calculus with bags in the asynchronous π -calculus. The general structure of the encoding is similar to our encoding of SILL₅ in the asynchronous π -calculus, modulo the richer syntax of SILL₅, which introduces more communication actions in the image of the encoding. We note that our encoding is greatly simplified by linearity and by the fact that SILL₅ does not employ *mixed* choice [31].

Linear Logic and Session Types. The propositions-as-types correspondence between linear logic and session types introduced by Caires and Pfenning [8, 9] initiated an ongoing line of research exploring the logical reading of sessions along various axes [42, 24, 34, 35, 2]. Starting with [8], which translates the linear session language into a π -calculus (which is more expressiveness than the source language), various works on encodings in this logical setting have been proposed [39, 41, 29, 28]. These study encodings between session-typed processes and functional languages, since the considered session languages are not powerful enough to express general π -calculus behaviors. Recent works [2, 12] attempt to address these limitations in expressiveness by allowing composed processes to share more than one linear channel, but still do not allow for the sharing available in SILL_5 , crucial to our encoding. We also highlight the work of Dardha and Pérez [14] comparing session-typed processes arising from linear logic and those from the Kobayashi-style typings [26, 25, 32] for the π -calculus. They observe that the *degree of sharing* determines an expressiveness hierarchy for typed processes and develop encodings from the latter into the former (not preserving the degree of sharing). In this sense, our encoding of asynchronous π -calculus completely preserves the sharing of channels, at the cost of allowing deadlocks when acquiring shared channels.

Session-Typed Behavioral Theory. The behavioral theory of session-typed processes has been studied in both the multiparty [27] and the linear logic settings [7, 34, 1]. Our notion of observation is related to the observed communication semantics of Atkey [1], which must also address the challenge of observing actions within a “closed-world” framework. However, their system is based on classical linear logic and does not have sharing, making the precise relationship with our formulation of observable on shared names unclear.

Substructural Logical Reasoning. The work of Deng et al. [15] studies a natural notion of logical preorder between linear logic contexts using process calculi techniques such as simulation preorders. While the study of the relationship between contexts can be seen as a study of multiset rewriting of configurations, the process calculus induced by their reading of linear logic is a fairly different formalism from SILL_5 . For instance, their labelled transition system cannot be reasonably used as a labelled transition system for SILL_5 since it cannot represent the equivalent of channel passing, nor does it make use of the deep inspection of multiset rewriting terms needed for our semantics and reasoning.

6 Concluding Remarks

In this paper, we gave an encoding of the untyped asynchronous π -calculus into SILL_5 via a universal shared session type \mathcal{U}_5 , proving its operational and observational correctness. This result shows that the full expressiveness of the untyped asynchronous π -calculus can be recovered in session-typed process calculi. We also provide an operationally correct encoding in the other direction to simulate shared session types in the π -calculus. Given their universality, session-typed calculi with manifest sharing become strong competitors over traditional approaches since they not only guarantee protocol compliance in the presence of non-determinism but also make sharing explicit in the type structure. For future work, we wish to investigate a general behavioral theory of manifest sharing, as well as study techniques to establish deadlock-freedom in the presence of shared channels.

References

- 1 Robert Atkey. Observed communication semantics for classical processes. In *European Symposium on Programming (ESOP)*, pages 56–82, 2017.
- 2 Robert Atkey, Sam Lindley, and J. Garrett Morris. Conflation confers concurrency. In S. Lindley et al., editor, *Wadler Festschrift*, pages 32–55. Springer LNCS 9600, 2016.
- 3 Stephanie Balzer and Frank Pfenning. Manifest sharing with session types. *Proceedings of the ACM on Programming Languages (PACMPL)*, 1(ICFP):37:1–37:29, 2017.
- 4 Stephanie Balzer and Frank Pfenning. Manifest sharing with session types. Technical Report CMU-CS-17-106, Carnegie Mellon University, March 2017.
- 5 Romain Beauxis, Catuscia Palamidessi, and Frank D. Valencia. On the asynchronous nature of the asynchronous π -calculus. In *Concurrency, Graphs and Models*, volume 5065 of *Lecture Notes in Computer Science*, pages 473–492. Springer, 2008.
- 6 P. N. Benton. A mixed linear and non-linear logic: Proofs, terms and models. In *8th International Workshop on Computer Science Logic (CSL)*, volume 933 of *Lecture Notes in Computer Science*, pages 121–135. Springer, 1994. An extended version appeared as Technical Report UCAM-CL-TR-352, University of Cambridge.
- 7 Luís Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. Behavioral polymorphism and parametricity in session-based communication. In *European Symposium on Programming (ESOP)*, pages 330–349. Springer, 2013.
- 8 Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *21st International Conference on Concurrency Theory (CONCUR)*, pages 222–236. Springer, 2010.
- 9 Luís Caires, Frank Pfenning, and Bernardo Toninho. Linear logic propositions as session types. *Mathematical Structures in Computer Science*, 26(3):367–423, 2016.
- 10 Iliano Cervesato and Andre Scedrov. Relating state-based and process-based concurrency through linear logic. *Information and Computation*, 207(10):1044–1077, 2009.
- 11 Karl Crary, Robert Harper, and Sidd Puri. What is a recursive module? In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 50–63, 1999.
- 12 Ornela Dardha and Simon J. Gay. A new linear logic for deadlock-free session-typed processes. In *Foundations of Software Science and Computation Structures (FoSSaCS)*, pages 91–109, 2018.
- 13 Ornela Dardha, Elena Giachino, and Davide Sangiorgi. Session types revisited. In *Principles and Practice of Declarative Programming (PPDP)*, pages 139–150, 2012.
- 14 Ornela Dardha and Jorge A. Pérez. Comparing deadlock-free session typed processes. In *EXPRESS/SOS*, pages 1–15, 2015.
- 15 Yuxin Deng, Robert J. Simmons, and Iliano Cervesato. Relating reasoning methodologies in linear logic and process algebra. *Mathematical Structure in Computer Science*, 26(5):868–906, 2016.
- 16 Simon J. Gay and Malcolm Hole. Subtyping for session types in the π -calculus. *Acta Informatica*, 42(2-3):191–225, 2005.
- 17 Hannah Gommerstadt, Limin Jia, and Frank Pfenning. Session-typed concurrent contracts. In A. Ahmed, editor, *European Symposium on Programming (ESOP'18)*, pages 771–798, Thessaloniki, Greece, 2018. Springer LNCS 10801.
- 18 Daniele Gorla. Towards a unified approach to encodability and separation results for process calculi. *Information and Computation*, 208(9):1031–1053, 2010.
- 19 Dennis Griffith and Frank Pfenning. SILL. <https://github.com/ISANobody/sill>, 2015.
- 20 Kohei Honda. Types for dyadic interaction. In *4th International Conference on Concurrency Theory (CONCUR)*, pages 509–523. Springer, 1993.

- 21 Kohei Honda and Mario Tokoro. An object calculus for asynchronous communication. In *5th European Conference on Object-Oriented Programming (ECOOP)*, Lecture Notes in Computer Science, pages 133–147. Springer, 1991.
- 22 Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In *7th European Symposium on Programming (ESOP)*, pages 122–138. Springer, 1998.
- 23 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 273–284. ACM, 2008.
- 24 Limin Jia, Hannah Gommerstadt, and Frank Pfenning. Monitors and blame assignment for higher-order session types. In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 582–594, 2016.
- 25 Naoki Kobayashi. A type system for lock-free processes. *Inf. Comput.*, 177(2):122–159, 2002.
- 26 Naoki Kobayashi. A new type system for deadlock-free processes. In *International Conference on Concurrency Theory (CONCUR)*, pages 233–247, 2006.
- 27 Dimitrios Kouzapas and Nobuko Yoshida. Globally governed session semantics. *Logical Methods in Computer Science*, 10(4), 2014.
- 28 Sam Lindley and J. Garrett Morris. A semantics for propositions as sessions. In *European Symposium On Programming (ESOP)*, pages 560–584, 2015.
- 29 Sam Lindley and J. Garrett Morris. Talking bananas: structural recursion for session types. In *International Colloquium on Functional Programming (ICFP)*, pages 434–447, 2016.
- 30 Robin Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999.
- 31 Uwe Nestmann. What is a "good" encoding of guarded choice? *Inf. Comput.*, 156(1-2):287–319, 2000.
- 32 Luca Padovani. Deadlock and lock freedom in the linear π -calculus. In *Computer Science Logic – Logic in Computer Science (CSL-LICS)*, pages 72:1–72:10, 2014.
- 33 Catuscia Palamidessi. Comparing the expressive power of the synchronous and asynchronous pi-calculi. *Mathematical Structures in Computer Science*, 13(5):685–719, 2003.
- 34 Jorge A. Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. Linear logical relations and observational equivalences for session-based concurrency. *Information and Computation*, 239:254–302, 2014.
- 35 Frank Pfenning and Dennis Griffith. Polarized substructural session types. In *18th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*, pages 3–22. Springer, 2015.
- 36 Jason Reed. A judgmental deconstruction of modal logic. Unpublished manuscript, January 2009. URL: <http://www.cs.cmu.edu/~jcreed/papers/jdml.pdf>.
- 37 Davide Sangiorgi and David Walker. *The π -Calculus - A Theory of Mobile Processes*. Cambridge University Press, 2001.
- 38 Bernardo Toninho. *A Logical Foundation for Session-based Concurrent Computation*. PhD thesis, Carnegie Mellon University and New University of Lisbon, 2015.
- 39 Bernardo Toninho, Luís Caires, and Frank Pfenning. Functions as session-typed processes. In *15th International Conference on Foundations of Software Science and Computational Structures (FOSSACS)*, pages 346–360. Springer, 2012.
- 40 Bernardo Toninho, Luís Caires, and Frank Pfenning. Higher-order processes, functions, and sessions: a monadic integration. In *22nd European Symposium on Programming (ESOP)*, pages 350–369. Springer, 2013.

30:18 A Universal Session Type for Untyped Asynchronous Communication

- 41 Bernardo Toninho and Nobuko Yoshida. On polymorphic sessions and functions - A tale of two (fully abstract) encodings. In *European Symposium On Programming (ESOP)*, pages 827–855, 2018.
- 42 Philip Wadler. Propositions as sessions. In *17th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 273–286. ACM, 2012.