Manifest Deadlock-Freedom for Shared Session Types

Stephanie Balzer[1(✉)], Bernardo Toninho[2(✉)], and Frank Pfenning[1]

[1] Carnegie Mellon University, Pittsburgh, USA
balzers@cs.cmu.edu
[2] NOVA LINCS, Universidade Nova de Lisboa, Lisbon, Portugal
btoninho@fct.unl.pt

Abstract. Shared session types generalize the Curry-Howard correspondence between intuitionistic linear logic and the session-typed π-calculus with adjoint modalities that mediate between linear and shared session types, giving rise to a programming model where shared channels must be used according to a locking discipline of acquire-release. While this generalization greatly increases the range of programs that can be written, the gain in expressiveness comes at the cost of deadlock-freedom, a property which holds for many linear session type systems. In this paper, we develop a type system for logically-shared sessions in which types capture not only the interactive behavior of processes but also constrain the order of resources (i.e., shared processes) they may acquire. This type-level information is then used to rule out cyclic dependencies among acquires and synchronization points, resulting in a system that ensures *deadlock-free communication* for well-typed processes in the presence of shared sessions, higher-order channel passing, and recursive processes. We illustrate our approach on a series of examples, showing that it rules out deadlocks in circular networks of both shared and linear recursive processes, while still being permissive enough to type concurrent implementations of shared imperative data structures as processes.

Keywords: Linear and shared session types · Deadlock-freedom

1 Introduction

Session types [25–27] naturally describe the interaction protocols that arise amongst concurrent processes that communicate via message-passing. This typing discipline has been integrated (with varying static safety guarantees) into several mainstream language such as Java [28,29], F# [43], Scala [49,50], Go [11] and Rust [33]. Session types moreover enjoy a logical correspondence between *linear logic* and the *session-typed π-calculus* [8,9,51,55]. Languages building on this correspondence [24,52,55] not only guarantee *session*

Supported by NSF Grant No. CCF-1718267: "Enriching Session Types for Practical Concurrent Programming" and NOVA LINCS (Ref. UID/CEC/04516/2019).

© The Author(s) 2019
L. Caires (Ed.): ESOP 2019, LNCS 11423, pp. 611–639, 2019.
https://doi.org/10.1007/978-3-030-17184-1_22

fidelity (i.e., type preservation) but also *deadlock-freedom* (i.e., global progress). The latter is guaranteed even in the presence of interleaved sessions, which are often excluded from the deadlock-free fragments of traditional session-typed frameworks [20,26,27,53]. These logical session types, however, exclude programming scenarios that demand *sharing* of mutable resources (e.g., shared databases or shared output devices) instead of functional resource replication.

To increase their practicality, logical session types have been extended with *manifest sharing* [2]. In the resulting language, linear and shared sessions coexist, but the type system enforces that clients of shared sessions run in mutual exclusion of each other. This separation is achieved by enforcing an *acquire-release* policy, where a client of a shared session must first acquire the session before it can participate in it along a private linear channel. Conversely, when a client releases a session, it gives up its linear channel and only retains a shared reference to the session. Thus, sessions in the presence of manifest sharing can change, or *shift*, between shared and linear execution modes. At the type-level, the acquire-release policy manifests in a stratification of session types into linear and shared with adjoint modalities [5,47,48], connecting the two strata. Operationally, the modality shifting *up* from the linear to the shared layer translates into an *acquire* and the one shifting *down* from shared to linear into a *release*.

Manifest sharing greatly increases the range of programs that can be written because it recovers the expressiveness of the untyped asynchronous π-calculus [3] while maintaining session fidelity. As in the π-calculus, however, the gain in expressiveness comes at the cost of *deadlock-freedom*. An illustrative example is an implementation of the classical dining philosophers problem, shown in Fig. 1, using the language $\mathsf{SILL_S}$ [2] that supports manifest sharing (in this setting we often equate a process with the session it offers along a distinguished channel). The code shows the process *fork_proc*, implementing a session of type sfork, and the processes *thinking* and *eating*, implementing sessions of type philosopher. We defer the details of the typing and the definition of the session types sfork and philosopher to Sect. 2 and focus on the programmatic working of the processes for now. For ease of reading, we typeset shared session types and variables denoting shared channel references in red.

A *fork_proc* process represents a fork that can be perpetually acquired and released. The actions accept and detach are the duals of acquire and release, respectively, allowing a process to accept an acquire by a client and to initiate a release by a client, respectively. Process thinking has two shared channel references as arguments, for the forks to the left and right of the philosopher, which the process tries to acquire. If the acquire succeeds, the process recurs as an eating philosopher with two (now) linear channel references of type lfork. Once a philosopher is done eating, it releases both forks and recurs as a thinking philosopher. Let's set a table for three philosopher that share three forks, all spawned as processes executing in parallel:

$$f_0 \leftarrow \mathit{fork_proc}\ ;\ f_1 \leftarrow \mathit{fork_proc}\ ;\ f_2 \leftarrow \mathit{fork_proc}\ ;$$
$$p_0 \leftarrow \mathit{thinking} \leftarrow f_0, f_1\ ;\ p_1 \leftarrow \mathit{thinking} \leftarrow f_1, f_2\ ;\ p_2 \leftarrow \mathit{thinking} \leftarrow f_2, f_0\ ;$$

```
fork_proc : {sfork}
c ← fork_proc =
  c' ← accept c ;
  c ← detach c' ;
  c ← fork_proc

thinking : {phil ← sfork, sfork}
c ← thinking ← left, right =
  left' ← acquire left ;
  right' ← acquire right ;
  c ← eating ← left', right' ;

eating : {phil ← lfork, lfork}
c ← eating ← left', right' =
  right ← release right' ;
  left ← release left' ;
  c ← thinking ← left, right
```

Fig. 1. Dining philosophers in $\mathsf{SILL_S}$ [2].

Infamously, this configuration may deadlock because of the *circular* dependency between the acquires. We can break this cycle by changing the last line to $p_2 \leftarrow$ *thinking* $\leftarrow f_0, f_2$, ensuring that forks are acquired in increasing order.

Perhaps surprisingly, cyclic dependencies between acquire requests are not the only source of deadlocks. Fig. 2 gives an example, defining the processes *owner* and *contester*, which both have a shared channel reference to a common resource that can be perpetually acquired and released. Both processes acquire the shared resource, but additionally exchange the message ping. More precisely, process *owner* spawns the process *contester*, acquires the shared resource, and only releases the resource after having received the message ping from the *contester*. Process *contester*, on the other hand, first attempts to acquire the resource and then sends the message *ping* to the owner. The program deadlocks if process *owner* acquires the resource first. In that case, process *owner* waits for process *contester* to send the message ping while process *contester* waits to acquire the resource held by process *owner*. We note that this deadlock arises in both synchronous and asynchronous semantics.

```
owner : {1 ← sres}
o ← owner ← sr =
  c ← contester ← sr ;
  lr ← acquire sr ;
  case c of
  | ping → wait c ;
          sr ← release lr ; close o

contester : {⊕{ping : 1} ← sres}
c ← contester ← sr =
  lr ← acquire sr ;
  c.ping ;
  sr ← release lr ;
  close c
```

Fig. 2. Circular dependencies among acquire and synchronization actions.

In this paper, we develop a type system for manifest sharing that rules out cycles between acquire requests and interdependencies between acquire requests and synchronization actions, detecting the two kinds of deadlocks explained above. In our type system, session types not only prescribe *when* resources must be acquired and released, but also the *range* of resources that may be acquired. To this end, we equip the type system with the notion of a *world*, an abstract value at which a process resides, and type processes relative to an acyclic *ordering* on worlds, akin to the partial-order based approaches of [34,37]. The contributions of this paper are:

- a characterization of the possible forms of deadlocks that can arise in shared session types;
- the introduction of manifest deadlock-freedom, where resource dependencies are manifest in the type structure via world modalities;
- its elaboration in the programming language $\mathsf{SILL}_{\mathsf{S}^+}$, resulting in a type system, a synchronous operational semantics, and proofs of session fidelity (preservation) and a strong form of progress that excludes all deadlocks;
- the novel abstraction of green and red arrows to reason about the interdependencies between processes;
- an illustration of the concepts on various examples, including an extensive comparison with related work.

This paper is structured as follows: Sect. 2 provides a short introduction to manifest sharing. Sect. 3 develops the type system and dynamics of the language $\mathsf{SILL}_{\mathsf{S}^+}$. Sect. 4 illustrates the introduced concepts on an extended example. Sect. 5 discusses the meta-theoretical properties of $\mathsf{SILL}_{\mathsf{S}^+}$, emphasizing progress. Sect. 6 compares with examples of related work and identifies future work. Sect. 7 discusses related work, and Sect. 8 concludes this paper.

2 Manifest Sharing

In the previous section, we have already explored the programmatic workings of *manifest sharing* [2], which enforces an *acquire-release* policy on shared channel references. In this section, we clarify the typing of shared processes.

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* [5,47,48], session types are partitioned into a linear and shared layer with two *adjoint modalities* connecting the layers:

$$\begin{array}{lcl} A_{\mathsf{S}} & \triangleq & \uparrow_{\mathsf{L}}^{\mathsf{S}} A_{\mathsf{L}} \\ A_{\mathsf{L}}, B_{\mathsf{L}} & \triangleq & A_{\mathsf{L}} \otimes B_{\mathsf{L}} \mid \oplus\{\overline{l : A_{\mathsf{L}}}\} \mid \&\{\overline{l : A_{\mathsf{L}}}\} \mid A_{\mathsf{L}} \multimap B_{\mathsf{L}} \mid \exists x{:}A_{\mathsf{S}}.B_{\mathsf{L}} \mid \Pi x{:}A_{\mathsf{S}}.B_{\mathsf{L}} \mid \mathbf{1} \mid \downarrow_{\mathsf{L}}^{\mathsf{S}} A_{\mathsf{S}} \end{array}$$

In the linear layer, we get the standard connectives of intuitionistic linear logic ($A_{\mathsf{L}} \otimes B_{\mathsf{L}}$, $A_{\mathsf{L}} \multimap B_{\mathsf{L}}$, $\oplus\{\overline{l : A_{\mathsf{L}}}\}$, $\&\{\overline{l : A_{\mathsf{L}}}\}$, and $\mathbf{1}$). These connectives are extended with the modal operator $\downarrow_{\mathsf{L}}^{\mathsf{S}} A_{\mathsf{S}}$, shifting *down* from the shared to the linear layer. Similarly, in the shared layer, we have the operator $\uparrow_{\mathsf{L}}^{\mathsf{S}} A_{\mathsf{L}}$, shifting *up* from the linear to the shared layer. The former translates into a *release* (and, dually, detach), the latter into an *acquire* (and, dually, accept). As a result, we obtain a system in which session types prescribe all forms of communication, including the acquisition and release of shared processes.

Table 1 provides an overview of $\mathsf{SILL}_{\mathsf{S}}$'s session types and their operational reading. Since $\mathsf{SILL}_{\mathsf{S}}$ is based on an intuitionistic interpretation of linear logic session types [8], types are expressed from the point of view of the *providing process* with the channel along which the process provides the session behavior being characterized by its session type. This choice avoids the explicit duality operation present in original presentations of session types [25,26] and in those based

Table 1. Session types in $\mathsf{SILL_S}$ and their operational meaning.

Session type		Process term		
current	**cont**	**current**	**cont**	**Description**
$c_{\mathsf{L}} : \oplus\{\overline{l : A_{\mathsf{L}}}\}$	$c_{\mathsf{L}} : A_{\mathsf{L}_h}$	$c_{\mathsf{L}}.l_h \; ; P$	P	sends label l_h along c_{L}
		$\mathsf{case}\ c_{\mathsf{L}}\ \mathsf{of}\ \overline{l \Rightarrow Q}$	Q_h	receives label l_h along c_{L}
$c_{\mathsf{L}} : \&\{\overline{l : A_{\mathsf{L}}}\}$	$c_{\mathsf{L}} : A_{\mathsf{L}_h}$	$\mathsf{case}\ c_{\mathsf{L}}\ \mathsf{of}\ \overline{l \Rightarrow P}$	P_h	receives label l_h along c
		$c_{\mathsf{L}}.l_h \; ; Q$	Q	sends label l_h along c_{L}
$c_{\mathsf{L}} : A_{\mathsf{L}} \otimes B_{\mathsf{L}}$	$c_{\mathsf{L}} : B_{\mathsf{L}}$	$\mathsf{send}\ c_{\mathsf{L}}\ d_{\mathsf{L}} \; ; P$	P	sends channel $d_{\mathsf{L}} : A_{\mathsf{L}}$ along c_{L}
		$y_{\mathsf{L}} \leftarrow \mathsf{recv}\ c_{\mathsf{L}} \; ; Q_{y_{\mathsf{L}}}$	$[d_{\mathsf{L}}/y_{\mathsf{L}}]\, Q_{y_{\mathsf{L}}}$	receives channel $d_{\mathsf{L}} : A_{\mathsf{L}}$ along c_{L}
$c_{\mathsf{L}} : A_{\mathsf{L}} \multimap B_{\mathsf{L}}$	$c_{\mathsf{L}} : B_{\mathsf{L}}$	$y_{\mathsf{L}} \leftarrow \mathsf{recv}\ c_{\mathsf{L}} \; ; P_{y_{\mathsf{L}}}$	$[d_{\mathsf{L}}/y_{\mathsf{L}}]\, P_{y_{\mathsf{L}}}$	receives channel $d_{\mathsf{L}} : A_{\mathsf{L}}$ along c_{L}
		$\mathsf{send}\ c_{\mathsf{L}}\ d_{\mathsf{L}} \; ; Q$	Q	sends channel $d_{\mathsf{L}} : A_{\mathsf{L}}$ along c_{L}
$c_{\mathsf{L}} : \Pi x{:}A_{\mathsf{S}}.B_{\mathsf{L}}$	$c_{\mathsf{L}} : B_{\mathsf{L}}$	$\mathsf{send}\ c_{\mathsf{L}}\ d_{\mathsf{S}} \; ; P$	P	sends channel $d_{\mathsf{S}} : A_{\mathsf{S}}$ along c_{L}
		$y_{\mathsf{S}} \leftarrow \mathsf{recv}\ c_{\mathsf{L}} \; ; Q_{y_{\mathsf{S}}}$	$[d_{\mathsf{S}}/y_{\mathsf{S}}]\, Q_{y_{\mathsf{S}}}$	receives channel $d_{\mathsf{S}} : A_{\mathsf{S}}$ along c_{L}
$c_{\mathsf{L}} : \exists x{:}A_{\mathsf{S}}.B_{\mathsf{L}}$	$c_{\mathsf{L}} : B_{\mathsf{L}}$	$y_{\mathsf{S}} \leftarrow \mathsf{recv}\ c_{\mathsf{L}} \; ; P_{y_{\mathsf{S}}}$	$[d_{\mathsf{S}}/y_{\mathsf{S}}]\, P_{y_{\mathsf{S}}}$	receives channel $d_{\mathsf{S}} : A_{\mathsf{S}}$ along c_{L}
		$\mathsf{send}\ c_{\mathsf{L}}\ d_{\mathsf{S}} \; ; Q$	Q	sends channel $d_{\mathsf{S}} : A_{\mathsf{S}}$ along c_{L}
$c_{\mathsf{L}} : \mathbf{1}$	-	$\mathsf{close}\ c_{\mathsf{L}}$	-	sends "end" along c_{L}
		$\mathsf{wait}\ c_{\mathsf{L}} \; ; Q$	Q	receives "end" along c_{L}
$c_{\mathsf{L}} : \downarrow_{\mathsf{L}}^{\mathsf{S}} A_{\mathsf{S}}$	$c_{\mathsf{S}} : A_{\mathsf{S}}$	$c_{\mathsf{S}} \leftarrow \mathsf{detach}\ c_{\mathsf{L}} \; ; P_{x_{\mathsf{S}}}$	$[c_{\mathsf{S}}/x_{\mathsf{S}}]\, P_{x_{\mathsf{S}}}$	sends "detach c_{S}" along c_{L}
		$x_{\mathsf{S}} \leftarrow \mathsf{release}\ c_{\mathsf{L}} \; ; Q_{x_{\mathsf{S}}}$	$[c_{\mathsf{S}}/x_{\mathsf{S}}]\, Q_{x_{\mathsf{S}}}$	receives "detach c_{S}" along c_{L}
$c_{\mathsf{S}} : \uparrow_{\mathsf{L}}^{\mathsf{S}} A_{\mathsf{L}}$	$c_{\mathsf{L}} : A_{\mathsf{L}}$	$c_{\mathsf{L}} \leftarrow \mathsf{acquire}\ c_{\mathsf{S}} \; ; Q_{x_{\mathsf{L}}}$	$[c_{\mathsf{L}}/x_{\mathsf{L}}]\, Q_{x_{\mathsf{L}}}$	sends "acquire c_{L}" along c_{S}
		$x_{\mathsf{L}} \leftarrow \mathsf{accept}\ c_{\mathsf{S}} \; ; P_{x_{\mathsf{L}}}$	$[c_{\mathsf{L}}/x_{\mathsf{L}}]\, P_{x_{\mathsf{L}}}$	receives "acquire c_{L}" along c_{S}

on classical linear logic [55]. Table 1 lists the points of view of the *provider* and *client* of a given connective in the first and second lines, respectively. Moreover, Table 1 gives for each connective its session type before and after the message exchange, along with their respective process terms. We can see that the process terms of a provider and a client for a given connective come in matching pairs, indicating that the participants' views of the session change consistently. We use the subscripts L and S to distinguish between linear and shared channels, respectively.

We are now able to give the session types of the processes *fork_proc*, *thinking*, and *eating* defined in the previous section:

$$\begin{aligned} &\mathsf{lfork} = \downarrow_{\mathsf{L}}^{\mathsf{S}} \mathsf{sfork} \\ &\mathsf{sfork} = \uparrow_{\mathsf{L}}^{\mathsf{S}} \mathsf{lfork} \\ &\mathsf{phil} = \mathbf{1} \end{aligned}$$

The mutually recursive session types lfork and sfork represent a fork that can perpetually be acquired and released. We adopt an *equi-recursive* [14] interpretation for recursive session types, silently equating a recursive type with its unfolding and requiring types to be *contractive* [19].

We briefly discuss the typing and the dynamics of acquire-release. The typing and the dynamics of the residual linear connectives are standard, and we detail them in the context of $\mathsf{SILL_{S^+}}$ (see Sect. 3). As is usual for an intuitionistic

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:

$$\text{(T-}\uparrow^{\mathsf{S}}_{\mathsf{L}}\text{R)}\ \frac{\Gamma; \cdot \vdash P_{x_{\mathsf{L}}} :: (x_{\mathsf{L}} : A_{\mathsf{L}})}{\Gamma \vdash x_{\mathsf{L}} \leftarrow \mathsf{accept}\ x_{\mathsf{S}}; P_{x_{\mathsf{L}}} :: (x_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}})} \qquad \text{(T-}\uparrow^{\mathsf{S}}_{\mathsf{L}}\text{L)}\ \frac{\Gamma, x_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}; \Delta, x_{\mathsf{L}} : A_{\mathsf{L}} \vdash Q_{x_{\mathsf{L}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}})}{\Gamma, x_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}; \Delta \vdash x_{\mathsf{L}} \leftarrow \mathsf{acquire}\ x_{\mathsf{S}}; Q_{x_{\mathsf{L}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}})}$$

$$\text{(T-}\downarrow^{\mathsf{S}}_{\mathsf{L}}\text{R)}\ \frac{\Gamma \vdash P_{x_{\mathsf{S}}} :: (x_{\mathsf{S}} : A_{\mathsf{S}})}{\Gamma; \cdot \vdash x_{\mathsf{S}} \leftarrow \mathsf{detach}\ x_{\mathsf{L}}; P_{x_{\mathsf{S}}} :: (x_{\mathsf{L}} : \downarrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{S}})} \qquad \text{(T-}\downarrow^{\mathsf{S}}_{\mathsf{L}}\text{L)}\ \frac{\Gamma, x_{\mathsf{S}} : A_{\mathsf{S}}; \Delta \vdash Q_{x_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}})}{\Gamma; \Delta, x_{\mathsf{L}} : \downarrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{S}} \vdash x_{\mathsf{S}} \leftarrow \mathsf{release}\ x_{\mathsf{L}}; Q_{x_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}})}$$

The typing judgments $\Gamma \vdash P :: (x_{\mathsf{S}} : A_{\mathsf{S}})$ and $\Gamma; \Delta \vdash P :: (x_{\mathsf{L}} : A_{\mathsf{L}})$ indicate that process P provides a session of type A along channel x, given the typing of the channels specified 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 precludes shared processes from depending on linear channel references [2,47], a restriction motivated from logic referred to as the independence principle [47]. Thus, when a shared session accepts an acquire and shifts to linear, it starts with an empty linear context.

Operationally, the dynamics of $\mathsf{SILL_S}$ is captured by *multiset rewriting rules* [12], 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 following:

(D-$\uparrow^{\mathsf{S}}_{\mathsf{L}}$)
$\mathsf{proc}(a_{\mathsf{S}},\ x_{\mathsf{L}} \leftarrow \mathsf{accept}\ a_{\mathsf{S}}\ ; P_{x_{\mathsf{L}}}),\ \mathsf{proc}(c_{\mathsf{L}},\ x_{\mathsf{L}} \leftarrow \mathsf{acquire}\ a_{\mathsf{S}}\ ; Q_{x_{\mathsf{L}}})$
$\longrightarrow\ \mathsf{proc}(a_{\mathsf{L}},\ [a_{\mathsf{L}}/x_{\mathsf{L}}]\ P_{x_{\mathsf{L}}}),\ \mathsf{proc}(c_{\mathsf{L}},\ [a_{\mathsf{L}}/x_{\mathsf{L}}]\ Q_{x_{\mathsf{L}}}),\ \mathsf{unavail}(a_{\mathsf{S}})$

(D-$\downarrow^{\mathsf{S}}_{\mathsf{L}}$)
$\mathsf{proc}(a_{\mathsf{L}},\ x_{\mathsf{S}} \leftarrow \mathsf{detach}\ a_{\mathsf{L}}\ ; P_{x_{\mathsf{S}}}),\ \mathsf{proc}(c_{\mathsf{L}},\ x_{\mathsf{S}} \leftarrow \mathsf{release}\ a_{\mathsf{L}}\ ; Q_{x_{\mathsf{S}}}),\ \mathsf{unavail}(a_{\mathsf{S}})$
$\longrightarrow\ \mathsf{proc}(a_{\mathsf{S}},\ [a_{\mathsf{S}}/x_{\mathsf{S}}]\ P_{x_{\mathsf{S}}}),\ \mathsf{proc}(c_{\mathsf{L}},\ [a_{\mathsf{S}}/x_{\mathsf{S}}]\ Q_{x_{\mathsf{S}}})$

Configuration states are defined by the predicates $\mathsf{proc}(c_m,\ P)$ and $\mathsf{unavail}(a_{\mathsf{S}})$. The former denotes a running process with process term P providing along channel c_m, the latter acts as 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 appear at one of two modes, a linear one, a_{L}, and a shared one, a_{S}. While the process (i.e. the session) 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. The channel at the appropriate mode is substituted for the variables occurring in process terms.

3 Manifest Deadlock-Freedom

In this section, we introduce our language $\mathsf{SILL_{S^+}}$, a session-typed language that supports sharing without deadlock. We focus on $\mathsf{SILL_{S^+}}$'s type system and dynamics in this section and discuss its meta-theoretical properties in Sect. 5.

3.1 Competition and Collaboration

The introduction of acquire-release, to ensure that the multiple clients of a shared process interact with the process in mutual exclusion from each other, gives rise to an obvious source of deadlocks, as acquire-release effectively amounts to a locking discipline. The typical approach to prevent deadlocks in that case is to impose a partial order on the resources and to *"lock-up"*, i.e., to lock the resources in ascending order. We adopted this strategy in Sect. 1 (Fig. 1) to break the cyclic dependencies among the acquires in the dining philosophers.

In Sect. 1, however, we also considered another example (Fig. 2) and discovered that *cyclic acquisitions* are not the only source of deadlocks, but deadlocks can also arise from *interdependent acquisitions and synchronizations*. In that example, we can prevent the deadlock by moving the acquire past the synchronization, in either of the two processes. Whereas in a purely linear session-typed system the sequencing of actions within a process do not affect other processes, the relative placement of acquire requests and synchronizations become relevant in a shared session-typed system.

Based on this observation, we can divide the processes in a shared-session discipline into *competitors* and *collaborators*. The former compete for a set of resources, whereas the latter do not overlap in the set of resources they acquire. For example, in the dining philosophers (Fig. 1), the philosophers p_0, p_1, and p_2 compete with each other for the set of forks f_0, f_1, and f_2, whereas the process that spawns the philosophers and the forks collaborates with either of them.

Transferring this idea to the process graph that emerges at run-time, we note that competitors are siblings whereas collaborators stand in a parent-descendant relationship. We illustrate this outcome on Fig. 3 that shows a possible run-time process graph for the dining philosophers. Linear processes are depicted as solid black circles with a white identifier and shared processes are depicted as dotted filled violet circles with a black identifier. Linear channels are depicted as black lines, shared channel references as dotted violet lines with the arrow head pointing to the shared process being acquired[1]. The identifiers P_0, P_1, and P_2 stand for the three philosophers, F_0, F_1, and F_2 for the three forks, and T for the process that sets the table. The current run-time graph depicts the scenario in which P_1 is eating, while the other two philosophers are still thinking.

Embedded in the graph is a *tree* that arises from the linear processes and the linear channels connecting them. For any two nodes in this tree, the *parent* node denotes the *client* process and the *child* node the *providing* process. We note that the *independence principle* (see Sect. 2), which precludes shared processes from depending on linear channel references, guarantees that there exists exactly one tree in the process graph, with the linear main process as its root. The shape of the tree changes when new processes are spawned, linear channels exchanged (through $\otimes$ and $\multimap$), or shared processes acquired. For example, process P_2 could acquire the shared fork F_0, which then becomes a linear child process of P_2, should the acquire succeed. As indicated by the shared channel references, the

[1] We have made sure to make the different concepts distinguishable in greyscale mode.

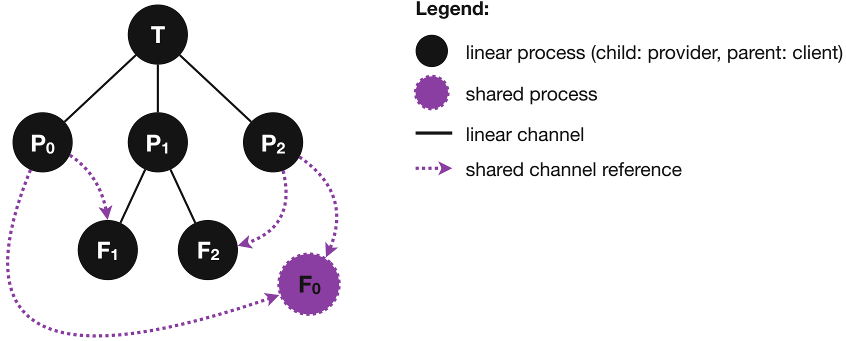

Fig. 3. Run-time process graph for dining philosophers (see Fig. 1).

sibling nodes P_0, P_1, and P_2 compete with each other for the nodes F_0, F_1, and F_2, whereas the node T does not compete for any of the resources acquired by its *descendants* (including F_1 and F_2). Our type system enforces this paradigm, as we discuss in the next section.

3.2 Type System

Invariants. Having identified the notions of *collaborators* and *competitors*, our type system must guarantee: *(i)* that collaborators acquire mutually disjoint sets of resources; *(ii)* that competitors employ a locking-up strategy for the resources they share; and, *(iii)* that competitors have released all acquired resources when synchronizing with other competitors. Invariant *(ii)* rules out cyclic acquisitions and invariants *(i)* and *(iii)* combined rule out interdependent acquisitions and synchronizations.

To express the high-level invariants above in our type system, we introduce the notion of a *world* – an abstract value that is equipped with a partial order – and associate such a world with every process. Programmers can *create* worlds, indicate the world at which a process resides at spawn time, and define an *order* on worlds. Moreover, we associate with each process a *range of worlds* that indicates the worlds of resources that the process may acquire. As a result, we obtain the following typing judgments:

$$\Psi;\, \Gamma \vdash P :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_k \updownarrow_{\omega_l}^{\omega_n}]) \quad \text{(where } \Psi^+ \text{ irreflexive)}$$

$$\Psi;\, \Gamma;\, \Phi;\, \Delta \vdash P :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_k \updownarrow_{\omega_l}^{\omega_n}]) \quad \text{(where } \Psi^+ \text{ irreflexive)}$$

The typing judgments reveal that we impose worlds at the *judgmental level*, resulting in a *hybrid system*, in which the adjoint modalities for acquire-release are complemented with world modalities that occur as *syntactic objects* in propositions [7]. We use the notation $x_m : A_m[\omega_k \updownarrow_{\omega_l}^{\omega_n}]$ (where m stands for S or L) to associate worlds ω_k, ω_l, and ω_n with a process that offers a session of type A_m along channel x. World ω_k denotes the world at which the process resides.

We refer to this world as the *self* world. Worlds ω_l and ω_n indicate the range of worlds of resources that the process may acquire, with ω_l denoting the *minimal (min)* world in this range and ω_n the *maximal (max)* one.

Process terms are typed relative to the order specified in Ψ and the contexts Γ, Φ, and Δ. As in Sect. 2, Γ is a structural context consisting of hypotheses on the typing of variables bound to shared channel references, augmented with world annotations. We find it necessary to split the linear context "Δ" from Sect. 2 into the two disjoint contexts Φ and Δ, allowing us to separate channels that are possibly aliased (due to sharing) from those that are not, respectively. Both Φ and Δ consist of hypotheses on the typing of variables that are bound to linear channels, augmented with world annotations. Ψ is presupposed to be *acyclic* and defined as: $\Psi \triangleq \cdot \mid \Psi',\ \omega_k < \omega_l \mid \Psi',\ \omega_o$, where ω stands for a concrete *world* w or a *world variable* δ. We allow Ψ to contain single worlds, to support singletons as well as to accommodate world creation prior to order declaration. We define the transitive closure Ψ^+, yielding a *strict partial order*, and the reflexive transitive closure Ψ^*, yielding a *partial order*.

The high-level invariants *(i)*, *(ii)*, and *(iii)* identified earlier naturally transcribe into the following invariants, which we impose on the typing judgments above. We use the notation $_\langle x_m \rangle; P$ to denote a process term that currently executes an action along channel x_m.

1. $\mathsf{min}(\mathsf{parent}) \leq \mathsf{self}(\mathsf{acquired_child}) \leq \mathsf{max}(\mathsf{parent})$:
   $\forall y_{\mathsf{L}} : B_{\mathsf{L}}[\omega_o \updownarrow^{\omega_r}_{\omega_p}] \in \Phi : \Psi^* \vdash \omega_l \leq \omega_o \leq \omega_n$
2. $\mathsf{max}(\mathsf{parent}) < \mathsf{min}(\mathsf{child})$:
   $\forall y_{\mathsf{L}} : B_{\mathsf{L}}[\omega_o \updownarrow^{\omega_r}_{\omega_p}] \in \Delta \cup \Phi : \Psi^+ \vdash \omega_n < \omega_p$
3. If $\Psi;\ \Gamma, x_{\mathsf{S}} : A[\omega_t \updownarrow^{\omega_v}_{\omega_u}];\ \Phi;\ \Delta \vdash x_{\mathsf{L}} \leftarrow \mathsf{acquire}\ x_{\mathsf{S}};\ Q_{x_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_k \updownarrow^{\omega_n}_{\omega_l}])$, then $\forall y_{\mathsf{L}} : B_{\mathsf{L}}[\omega_o \updownarrow^{\omega_r}_{\omega_p}] \in \Phi : \Psi^+ \vdash \omega_o < \omega_t$.
4. If $\Psi;\ \Gamma;\ \Phi;\ \Delta \vdash _\langle x_m \rangle; P :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_k \updownarrow^{\omega_n}_{\omega_l}])$, then $\Phi = (\cdot)$.

Invariants 1 and 2 ensure that, for any node in the tree, the acquired resources reside at smaller worlds than those acquired by any descendant. As a result, the two invariants guarantee high-level invariant *(i)*. Invariant 3, on the other hand, imposes a lock-up strategy on acquires and thus guarantees high-level invariant *(ii)*. To guarantee high-level invariant *(iii)*, we impose Invariant 4, which forces a process to release any acquired resources before communicating along its offering channel. Since sibling nodes cannot be directly connected by a linear channel, the only way for them to synchronize is through a common parent. Finally, to guarantee that world annotations are internally consistent, we require for each annotation $[\omega_k \updownarrow^{\omega_n}_{\omega_l}]$ that $\omega_k < \omega_l \leq \omega_n$.

Rules. We now present select process typing rules, a complete listing is provided in the companion technical report [4]. The only new rules with respect to the language $\mathsf{SILL_S}$ [2] are those pertaining to world creation and order determination. These are extra-logical judgmental rules. We allow both linear and shared processes to create and relate worlds. Rules (T-New$_{\mathsf{L}}$) and (T-New$_{\mathsf{S}}$) create a new world w and make it available to the continuation Q_{w}. Rules (T-Ord$_{\mathsf{L}}$) and (T-Ord$_{\mathsf{S}}$) relate two existing worlds, while preserving acyclicity of the order.

$$\frac{\Psi, \mathsf{w};\ \Gamma;\ \Phi;\ \Delta \vdash Q_{\mathsf{w}} :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma;\ \Phi;\ \Delta \vdash \mathsf{w} \leftarrow \mathsf{new_world};\ Q_{\mathsf{w}} :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-New}_{\mathsf{L}})$$

$$\frac{\Psi, \mathsf{w};\ \Gamma \vdash Q_{\mathsf{w}} :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma \vdash \mathsf{w} \leftarrow \mathsf{new_world};\ Q_{\mathsf{w}} :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-New}_{\mathsf{S}})$$

$$\frac{\omega_p, \omega_r \in \Psi \qquad (\Psi, \omega_p < \omega_r)^+ \text{ irreflexive} \quad \Psi, \omega_p < \omega_r;\ \Gamma;\ \Phi;\ \Delta \vdash Q :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma;\ \Phi;\ \Delta \vdash \omega_p < \omega_r;\ Q :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-Ord}_{\mathsf{L}})$$

$$\frac{\omega_p, \omega_r \in \Psi \qquad (\Psi, \omega_p < \omega_r)^+ \text{ irreflexive} \quad \Psi, \omega_p < \omega_r;\ \Gamma \vdash Q :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma \vdash \omega_p < \omega_r;\ Q :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-Ord}_{\mathsf{S}})$$

We now consider the typing rule for acquire, which must explicitly enforce the various low-level invariants above. Since an acquire results in the addition of a new child node to the executing process, the rule can interfere with Invariants 1 and 2. The first two premises of the rule ensure that the two invariants are preserved. Moreover, the rule has to ensure that the acquiring process is locking-up (Invariant 3), which is achieved by the third premise.

$$\frac{\Psi^* \vdash \omega_k \leq \omega_m \leq \omega_n \qquad \Psi^+ \vdash \omega_n < \omega_u \qquad \forall y_{\mathsf{L}} : B_{\mathsf{L}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}] \in \Phi : \omega_l < \omega_m \quad \Psi;\ \Gamma, x_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Phi, x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash Q_{x_{\mathsf{L}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma, x_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Phi;\ \Delta \vdash x_{\mathsf{L}} \leftarrow \mathsf{acquire}\, x_{\mathsf{S}}\, ;Q_{x_{\mathsf{L}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\uparrow^{\mathsf{S}}_{\mathsf{L}}\text{L})$$

The remaining shift rules are actually *unchanged* with respect to $\mathsf{SILL_S}$, modulo the world annotations. In particular, low-level Invariant 4 is already satisfied because the conclusion of rule (T-$\uparrow^{\mathsf{S}}_{\mathsf{L}}$R) does not have a context Φ and because the independence principle forces Φ to be empty in rule (T-$\downarrow^{\mathsf{S}}_{\mathsf{L}}$R).

$$\frac{\Psi;\ \Gamma;\ \cdot;\ \cdot \vdash P_{x_{\mathsf{L}}} :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma \vdash x_{\mathsf{L}} \leftarrow \mathsf{accept}\, x_{\mathsf{S}}\, ;P_{x_{\mathsf{L}}} :: (x_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-}\uparrow^{\mathsf{S}}_{\mathsf{L}}\text{R})$$

$$\frac{\Psi;\ \Gamma, x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Phi;\ \Delta \vdash Q_{x_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\ \Phi, x_{\mathsf{L}} : \downarrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash x_{\mathsf{S}} \leftarrow \mathsf{release}\, x_{\mathsf{L}}\, ;Q_{x_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\downarrow^{\mathsf{S}}_{\mathsf{L}}\text{L})$$

$$\frac{\Psi;\ \Gamma \vdash P_{x_{\mathsf{S}}} :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma;\ \cdot;\ \cdot \vdash x_{\mathsf{S}} \leftarrow \mathsf{detach}\, x_{\mathsf{L}}\, ;P_{x_{\mathsf{S}}} :: (x_{\mathsf{L}} : \downarrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{S}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-}\downarrow^{\mathsf{S}}_{\mathsf{L}}\text{R})$$

We now consider the linear connectives, starting with $\mathbf{1}$. Rule (T-$\mathbf{1}_{\mathrm{L}}$) reveals that only processes that have never been acquired may be terminated. This restriction is important to guarantee progress because existing clients of a shared process may wait indefinitely otherwise. We impose the restriction as a well-formedness condition on a session type, giving rise to a *strictly equi-synchronizing* session type. The notion of an *equi-synchronizing* session type [2] has been defined for $\mathsf{SILL_S}$ and guarantees that a process that has been acquired at a type A_{S} is released back to the type A_{S}, should it ever be released. A *strictly* equi-synchronizing session type additionally requires that an acquired resource *must* be released. The corresponding rules can be found in [4]. Linearity enforces Invariant 4 in rule (T-$\mathbf{1}_{\mathrm{R}}$), making sure that no linear channels are left behind.

$$\frac{\Psi;\ \Gamma;\ \Phi;\ \Delta \vdash Q :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\ \Phi;\ \Delta, x_{\mathsf{L}} : \mathbf{1}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash \mathsf{wait}\, x_{\mathsf{L}}\, ; Q :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\mathbf{1}_{\text{L}})$$

$$\frac{}{\Psi;\ \Gamma;\ \cdot;\ \cdot \vdash \mathsf{close}\, x_{\mathsf{L}} :: (x_{\mathsf{L}} : \mathbf{1}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-}\mathbf{1}_{\text{R}})$$

Next, we consider internal and external choice. Since internal and external choice cannot alter the linear process tree of a process graph, the rules are very similar to the ones in $\mathsf{SILL_S}$. The only differences are that we get two left rules for each connective and that the Φ-context of each right rule must be empty to satisfy Invariant 4. The former is merely due to the tracking of possibly aliased sessions in the Φ context. We only list rules for internal choice, those for external choice are dual and can be found in [4].

$$\frac{(\forall i)\ \ \Psi;\ \Gamma;\ \Phi;\ \Delta, x_{\mathsf{L}} : A_{\mathsf{L}_i}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash Q_i :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\ \Phi;\ \Delta, x_{\mathsf{L}} : \oplus\{\overline{l : A_{\mathsf{L}}}\}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash \mathsf{case}\ x_{\mathsf{L}}\ \mathsf{of}\ \overline{l \Rightarrow Q} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\oplus_{\text{L}_1})$$

$$\frac{(\forall i)\ \ \Psi;\ \Gamma;\ \Phi, x_{\mathsf{L}} : A_{\mathsf{L}_i}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash Q_i :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\ \Phi, x_{\mathsf{L}} : \oplus\{\overline{l : A_{\mathsf{L}}}\}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash \mathsf{case}\ x_{\mathsf{L}}\ \mathsf{of}\ \overline{l \Rightarrow Q} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\oplus_{\text{L}_2})$$

$$\frac{\Psi;\ \Gamma;\ \cdot;\ \Delta \vdash P :: (x_{\mathsf{L}} : A_{\mathsf{L}\,h}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma;\ \cdot;\ \Delta \vdash x_{\mathsf{L}}.l_h\, ; P :: (x_{\mathsf{L}} : \oplus\{\overline{l : A_{\mathsf{L}}}\}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-}\oplus_{\text{R}})$$

More interesting are linear channel output and input, since these alter the linear process tree of a process graph. Moreover, additional world annotations are needed to indicate the worlds of the channel that is exchanged. For the latter we use the notation $@\omega_l \updownarrow^{\omega_r}_{\omega_p}$, indicating that the exchanged channel has the worlds ω_l, ω_p, and ω_r for self, min, and max, respectively. To account for induced changes in the process graph, the rules that type an input of a linear channel must guard against any disturbance of Invariants 1 and 2. Because the two invariants guarantee that parents do not overlap with their descendants in terms of acquired resources, they prevent any exchange of acquired channels. We thus restrict $\otimes$ and $\multimap$ to the exchange of channels that have not yet been acquired. This is not a limitation since, as we will see below, shared channel output and input are unrestricted.

Even with the above restriction in place, we still have to make sure that a received channel satisfies Invariant 2. If we were to state a corresponding premise on the receiving rules, invertibility of the rules would be disturbed. To uphold invertibility, we impose a well-formedness condition on session types that ensures for a session of type $A_{\mathsf{L}}@\omega_l \updownarrow^{\omega_r}_{\omega_p} \otimes B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}]$ that $\omega_v < \omega_p$ and, analogously, for a session of type $A_{\mathsf{L}}@\omega_l \updownarrow^{\omega_r}_{\omega_p} \multimap B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}]$ that $\omega_v < \omega_p$. Session types are checked to be well-formed upon process definition. Given type well-formedness, we obtain the following rules for $\multimap$, noting that the right rule enforces Invariant 4 by requiring an empty Φ-context. The rules for $\otimes$ are dual.

$$\frac{\Psi;\ \Gamma;\ \Phi;\ \Delta, x_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash Q :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\ \Phi;\ \Delta, x_{\mathsf{L}} : A_{\mathsf{L}}@\omega_l \updownarrow^{\omega_r}_{\omega_p} \multimap B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}], y_{\mathsf{L}} : A_{\mathsf{L}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}] \vdash \mathsf{send}\ x_{\mathsf{L}}\ y_{\mathsf{L}}\ ; Q :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}{\multimap}\text{L}_1)$$

$$\frac{\Psi;\ \Gamma;\ \Phi, x_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash Q :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\ \Phi, x_{\mathsf{L}} : A_{\mathsf{L}}@\omega_l \updownarrow^{\omega_r}_{\omega_p} \multimap B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta, y_{\mathsf{L}} : A_{\mathsf{L}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}] \vdash \mathsf{send}\ x_{\mathsf{L}}\ y_{\mathsf{L}}\ ; Q :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}{\multimap}\text{L}_2)$$

$$\frac{\Psi;\ \Gamma;\ \cdot\ ;\ \Delta, y_{\mathsf{L}} : A_{\mathsf{L}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}] \vdash P_{y_{\mathsf{L}}} :: (x_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma;\ \cdot\ ;\ \Delta \vdash y_{\mathsf{L}} \leftarrow \mathsf{recv}\ x_{\mathsf{L}}\ ; P_{y_{\mathsf{L}}} :: (x_{\mathsf{L}} : A_{\mathsf{L}}@\omega_l \updownarrow^{\omega_r}_{\omega_p} \multimap B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-}{\multimap}\text{R})$$

Since there are no invariants imposed on the shared context Γ, the rules for shared channel output and input are identical to those in $\mathsf{SILL_S}$. The only differences are that we have two left rules and that the Φ-context of the right rule must be empty to satisfy Invariant 4. The former is merely due to the tracking of possibly aliased sessions in the Φ context.

$$\frac{\Psi;\ \Gamma, y_{\mathsf{S}} : A_{\mathsf{S}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}];\Phi;\ \Delta, x_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash Q_{y_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\Phi;\ \Delta, x_{\mathsf{L}} : \exists x{:}A_{\mathsf{S}}@\omega_l \updownarrow^{\omega_r}_{\omega_p}\ .\ B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash y_{\mathsf{S}} \leftarrow \mathsf{recv}\ x_{\mathsf{L}}\ ; Q_{y_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\exists\text{L}_1)$$

$$\frac{\Psi;\ \Gamma, y_{\mathsf{S}} : A_{\mathsf{S}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}];\Phi, x_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash Q_{y_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}{\Psi;\ \Gamma;\Phi, x_{\mathsf{L}} : \exists x{:}A_{\mathsf{S}}@\omega_l \updownarrow^{\omega_r}_{\omega_p}\ .\ B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}];\ \Delta \vdash y_{\mathsf{S}} \leftarrow \mathsf{recv}\ x_{\mathsf{L}}\ ; Q_{y_{\mathsf{S}}} :: (z_{\mathsf{L}} : C_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-}\exists\text{L}_2)$$

$$\frac{\Psi;\ \Gamma, y_{\mathsf{S}} : A_{\mathsf{S}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}];\cdot\ ;\ \Delta \vdash P :: (x_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}{\Psi;\ \Gamma, y_{\mathsf{S}} : A_{\mathsf{S}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}];\cdot\ ;\ \Delta \vdash \mathsf{send}\ x_{\mathsf{L}}\ y_{\mathsf{S}}\ ; P :: (x_{\mathsf{L}} : \exists x{:}A_{\mathsf{S}}@\omega_l \updownarrow^{\omega_r}_{\omega_p}\ .\ B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}])}\ (\text{T-}\exists\text{R})$$

We finally consider the rules for forwarding and spawning. We allow a shared forward between processes that offer the same session at the same worlds. Because forwards have to be *world-invariant*, however, no well-typed program could ever have a linear forward. The process being forwarded to must be in either of the contexts Φ or Δ, and thus satisfies Invariant 2, making it impossible for the world annotations of the forwarder and forwardee to match. We omit linear forwarding and discuss possible future extensions in Sect. 6.

$$\frac{}{\Psi;\ \Gamma, y_{\mathsf{S}} : A_{\mathsf{S}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \vdash \mathsf{fwd}\ x_{\mathsf{S}}\ y_{\mathsf{S}} :: (x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-Id}_{\mathsf{S}})$$

The rules for spawning depend on the possible modes of the spawning and spawned processes: (T-Spawn$_{\mathsf{LL}}$) specifies how a linear process can spawn another linear process; (T-Spawn$_{\mathsf{SS}}$) specifies how a shared processes can spawn another shared process. The rules are checked relative to a process definition found in the signature Σ and to a world substitution mapping $\gamma : |\Psi| \rightarrow |\Psi'|$, such that for each $\delta \in \Psi'$ we have $\Psi \vdash \gamma(\delta)$, where $|\Psi|$ denotes the *field* of Ψ (i.e., the union of its domain and range). As usual, we lift substitution to types $\hat{\gamma}(A_m)$, contexts $\hat{\gamma}(\Gamma)$, and orders $\hat{\gamma}(\Psi)$. Both rules ensure that, given the mapping γ, the order Ψ of the spawning process entails the one of the process definition ($\Psi \vdash \hat{\gamma}(\Psi')$). The linear spawn rule (T-Spawn$_{\mathsf{LL}}$) further enforces Invariant 2 for the spawned child. We note that the spawned child enters the linear context Δ in the spawning process' continuation since no aliases to such a process can exist at this point.

$$\dfrac{\begin{array}{c}\Delta_1 = \overline{y_{\mathsf{L}} : B_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}]} \qquad \Phi_1 = \overline{\tilde{y}_{\mathsf{L}} : \tilde{B}_{\mathsf{L}}[\tilde{\omega}_m \updownarrow^{\tilde{\omega}_v}_{\tilde{\omega}_u}]} \qquad \Gamma_1 = \overline{z_{\mathsf{S}} : C_{\mathsf{S}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}]} \\ (\Psi' \vdash x'_{\mathsf{L}} : A'_{\mathsf{L}}[\delta_j \updownarrow^{\delta_n}_{\delta_k}] \leftarrow X_{\mathsf{L}} \leftarrow \Delta', \Phi', \Gamma' = P_{x'_{\mathsf{L}}, \mathsf{dom}(\Delta'), \mathsf{dom}(\Phi'), \mathsf{dom}(\Gamma'), \Psi''}) \in \Sigma \\ \hat{\gamma}(A'_{\mathsf{L}}[\delta_j \updownarrow^{\delta_n}_{\delta_k}]) = A_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \qquad \hat{\gamma}(\Delta') = \Delta_1 \qquad \hat{\gamma}(\Phi') = \Phi_1 \qquad \hat{\gamma}(\Gamma') = \Gamma_1 \qquad \Psi \vdash \hat{\gamma}(\Psi') \\ \Psi^+ \vdash \omega_t < \omega_k \\ \Psi;\ \Gamma_1, \Gamma_2;\ \Phi_2;\ \Delta_2, x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \vdash Q_{x_{\mathsf{L}}} :: (z''_{\mathsf{L}} : D_{\mathsf{L}}[\omega_i \updownarrow^{\omega_t}_{\omega_q}])\end{array}}{\Psi;\ \Gamma_1, \Gamma_2;\ \Phi_1, \Phi_2;\ \Delta_1, \Delta_2 \vdash x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \leftarrow X_{\mathsf{L}} \leftarrow \overline{y_{\mathsf{L}}}, \overline{\tilde{y}_{\mathsf{L}}}, \overline{z_{\mathsf{S}}};\ Q_{x_{\mathsf{L}}} :: (z''_{\mathsf{L}} : D_{\mathsf{L}}[\omega_i \updownarrow^{\omega_t}_{\omega_q}])} \ (\text{T-Spawn}_{\mathsf{LL}})$$

$$\dfrac{\begin{array}{c}\Gamma_1 = \overline{z_{\mathsf{S}} : C_{\mathsf{S}}[\omega_l \updownarrow^{\omega_r}_{\omega_p}]} \qquad (\Psi' \vdash x'_{\mathsf{S}} : A'_{\mathsf{S}}[\delta_j \updownarrow^{\delta_n}_{\delta_k}] \leftarrow X_{\mathsf{S}} \leftarrow \Gamma' = P_{x'_{\mathsf{S}}, \mathsf{dom}(\Gamma'), \Psi''}) \in \Sigma \\ \hat{\gamma}(A'_{\mathsf{S}}[\delta_j \updownarrow^{\delta_n}_{\delta_k}]) = A_{\mathsf{S}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \qquad \hat{\gamma}(\Gamma') = \Gamma_1 \qquad \Psi \vdash \hat{\gamma}(\Psi') \\ \Psi;\ \Gamma_1, \Gamma_2, x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \vdash Q_{x_{\mathsf{S}}} :: (z''_{\mathsf{S}} : D_{\mathsf{S}}[\omega_i \updownarrow^{\omega_t}_{\omega_q}])\end{array}}{\Psi;\ \Gamma_1, \Gamma_2 \vdash x_{\mathsf{S}} : A_{\mathsf{S}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}] \leftarrow X_{\mathsf{S}} \leftarrow \overline{z_{\mathsf{S}}};\ Q_{x_{\mathsf{S}}} :: (z''_{\mathsf{S}} : D_{\mathsf{S}}[\omega_i \updownarrow^{\omega_t}_{\omega_q}])} \ (\text{T-Spawn}_{\mathsf{SS}})$$

In the companion technical report [4], we provide a variant of rule (T-Spawn$_{\mathsf{LL}}$) for the case of a linear recursive tail call. Without linear forwarding, a linear tail call can no longer be implicitly "de-sugared" into a spawn and a linear forward [2,22,52], but must be accounted for explicitly. In the report, we also provide the rules for checking process definitions. Those rules make sure that the process' world order is acyclic, that the types of the providing session and argument sessions are well-formed, and that the process satisfies Invariants 1 and 2.

3.3 Dining Philosophers in $\mathsf{SILL}_{\mathsf{S}^+}$

Having introduced our type system, we revisit the dining philosophers from Sect. 1 and show how to program the example in $\mathsf{SILL}_{\mathsf{S}^+}$, ensuring that the program will run without deadlocks. The code is given in Fig. 4. We note the world annotations in the signature of the process definitions. For instance,

$$\mathit{thinking} : \{\delta_0 < \delta_1, \delta_1 < \delta_2, \delta_2 < \delta_3 \vdash \mathsf{phil}[\delta_0 \updownarrow^{\delta_2}_{\delta_1}] \leftarrow \mathsf{sfork}[\delta_1 \updownarrow^{\delta_3}_{\delta_3}], \mathsf{sfork}[\delta_2 \updownarrow^{\delta_3}_{\delta_3}]; \cdot; \cdot\}$$

indicates that, given the order $\delta_0 < \delta_1 < \delta_2 < \delta_3$, process *thinking* provides a session of type $\mathsf{phil}[\delta_0 \updownarrow^{\delta_2}_{\delta_1}]$ and uses two shared channel references of type $\mathsf{sfork}[\delta_1 \updownarrow^{\delta_3}_{\delta_3}]$ and $\mathsf{sfork}[\delta_2 \updownarrow^{\delta_3}_{\delta_3}]$. The two $\cdot$ signify that neither acquired nor linear channel references are given as arguments. The signature indicates that the two shared fork references reside at different worlds, such that the world of the first one is smaller than the one of the second.

Let's briefly convince ourselves that the two acquires in process *thinking* in Fig. 4 are type-correct. For each acquire we have to show that: the world of the resource to be acquired is within the acquiring process' range; the max of the acquiring process is smaller than the min of the acquired resource; and, that the self of the acquired resource is larger than those of all already acquired resources. We can convince ourselves that all those conditions are readily met.

$$\begin{array}{l}
\textit{thinking} : \{\delta_0 < \delta_1, \delta_1 < \delta_2, \delta_2 < \delta_3 \vdash \\
\quad \mathsf{phil}[\delta_0 \updownarrow_{\delta_1}^{\delta_2}] \leftarrow \mathsf{sfork}[\delta_1 \updownarrow_{\delta_3}^{\delta_3}], \mathsf{sfork}[\delta_2 \updownarrow_{\delta_3}^{\delta_3}]; \cdot; \cdot\} \\
c[\delta_0 \updownarrow_{\delta_1}^{\delta_2}] \leftarrow \textit{thinking} \leftarrow \textit{left}[\delta_1 \updownarrow_{\delta_3}^{\delta_3}], \textit{right}[\delta_2 \updownarrow_{\delta_3}^{\delta_3}] = \\
\quad \textit{left}' \leftarrow \mathsf{acquire}\ \textit{left}\ ; \\
\quad \textit{right}' \leftarrow \mathsf{acquire}\ \textit{right}\ ; \\
\quad c \leftarrow \textit{eating} \leftarrow \textit{left}', \textit{right}'\ ;
\end{array}
\qquad
\begin{array}{l}
\mathsf{lfork} = \downarrow_{\mathsf{L}}^{\mathsf{S}} \mathsf{sfork} \\
\mathsf{sfork} = \uparrow_{\mathsf{L}}^{\mathsf{S}} \mathsf{lfork} \\
\mathsf{phil} = \mathbf{1}
\end{array}$$

$$\begin{array}{l}
\textit{eating} : \{\delta_0 < \delta_1, \delta_1 < \delta_2, \delta_2 < \delta_3 \vdash \\
\quad \mathsf{phil}[\delta_0 \updownarrow_{\delta_1}^{\delta_2}] \leftarrow \cdot; \mathsf{lfork}[\delta_1 \updownarrow_{\delta_3}^{\delta_3}], \mathsf{lfork}[\delta_2 \updownarrow_{\delta_3}^{\delta_3}]; \cdot\} \\
c[\delta_0 \updownarrow_{\delta_1}^{\delta_2}] \leftarrow \textit{eating} \leftarrow \textit{left}'[\delta_1 \updownarrow_{\delta_3}^{\delta_3}], \textit{right}'[\delta_2 \updownarrow_{\delta_3}^{\delta_3}] = \\
\quad \textit{right} \leftarrow \mathsf{release}\ \textit{right}'\ ; \\
\quad \textit{left} \leftarrow \mathsf{release}\ \textit{left}'\ ; \\
\quad c \leftarrow \textit{thinking} \leftarrow \textit{left}, \textit{right}
\end{array}
\qquad
\begin{array}{l}
\textit{fork_proc} : \{\delta_0 < \delta_1 \vdash \mathsf{sfork}[\delta_0 \updownarrow_{\delta_1}^{\delta_1}]\} \\
c[\delta_0 \updownarrow_{\delta_1}^{\delta_1}] \leftarrow \textit{fork_proc} = \\
\quad c' \leftarrow \mathsf{accept}\ c\ ; \\
\quad c \leftarrow \mathsf{detach}\ c'\ ; \\
\quad c'' : \mathsf{sfork}[\delta_0 \updownarrow_{\delta_1}^{\delta_1}] \leftarrow \textit{fork_proc}\ ; \\
\quad \mathsf{fwd}\ c\ c''
\end{array}$$

Fig. 4. Deadlock-free version of dining philosophers in $\mathsf{SILL}_{\mathsf{S}+}$.

We note, however, that if we were to swap the two acquires, the program would not type-check.

Let us once more set the table for three philosophers and three forks. We execute this code in a process with world annotations $[\delta_a \updownarrow_{\delta_b}^{\delta_b}]$ such that $\delta_a < \delta_b$. We first create new worlds and define their order:

$$\begin{array}{l}
\mathsf{w}_1 \leftarrow \mathsf{new_world}; \mathsf{w}_2 \leftarrow \mathsf{new_world}; \mathsf{w}_3 \leftarrow \mathsf{new_world}; \mathsf{w}_4 \leftarrow \mathsf{new_world}; \\
\delta_a < \mathsf{w}_1; \delta_a < \mathsf{w}_2; \delta_b < \mathsf{w}_1; \mathsf{w}_1 < \mathsf{w}_2; \mathsf{w}_1 < \mathsf{w}_3; \mathsf{w}_1 < \mathsf{w}_4; \mathsf{w}_2 < \mathsf{w}_3; \mathsf{w}_2 < \mathsf{w}_4; \mathsf{w}_3 < \mathsf{w}_4;
\end{array}$$

We then spawn the forks, each residing at a different world, such that the max world of a fork is higher than the self of the highest fork, ensuring Invariant 2 for the philosopher processes that we spawn afterwards:

$$\begin{array}{l}
f_1 : \mathsf{sfork}[\mathsf{w}_1 \updownarrow_{\mathsf{w}_4}^{\mathsf{w}_4}] \leftarrow \textit{fork_proc}\ ;\ f_2 : \mathsf{sfork}[\mathsf{w}_2 \updownarrow_{\mathsf{w}_4}^{\mathsf{w}_4}] \leftarrow \textit{fork_proc}\ ; \\
f_3 : \mathsf{sfork}[\mathsf{w}_3 \updownarrow_{\mathsf{w}_4}^{\mathsf{w}_4}] \leftarrow \textit{fork_proc}\ ;
\end{array}$$

When we spawn the philosophers, we ensure that P_0 is going to pick up fork F_1 and then F_2, P_1 is going to pick up F_2 and then F_3, and P_2 is going to pick up F_1 and then F_3.

$$\begin{array}{l}
p_0 : \mathsf{phil}[\delta_a \updownarrow_{\mathsf{w}_1}^{\mathsf{w}_2}] \leftarrow \textit{thinking} \leftarrow \cdot; \cdot; f_1, f_2\ ;\ p_1 : \mathsf{phil}[\delta_a \updownarrow_{\mathsf{w}_2}^{\mathsf{w}_3}] \leftarrow \textit{thinking} \leftarrow \cdot; \cdot; f_2, f_3\ ; \\
p_2 : \mathsf{phil}[\delta_a \updownarrow_{\mathsf{w}_1}^{\mathsf{w}_3}] \leftarrow \textit{thinking} \leftarrow \cdot; \cdot; f_1, f_3\ ;
\end{array}$$

We note that the deadlocking spawn

$$p_2 : \mathsf{phil}[\delta_a \updownarrow_{\mathsf{w}_1}^{\mathsf{w}_3}] \leftarrow \textit{thinking} \leftarrow \cdot; \cdot; f_3, f_1\ ;$$

is type-incorrect since we would substitute both w_1 and w_3 for δ_1 and w_3 and w_1 for δ_2, which violates the ordering constraints put in place by typing.

3.4 Dynamics

We now give the *dynamics* of $\mathsf{SILL}_{\mathsf{S}^+}$. Our current system is based on a *synchronous* dynamics. While this choice is more conservative, it allows us to narrow the complexity of the problem at hand.

As in $\mathsf{SILL}_{\mathsf{S}}$, we use *multiset rewriting rules* [12] to capture the dynamics of $\mathsf{SILL}_{\mathsf{S}^+}$ (see Sect. 2). Multiset rewriting rules represent computation in terms of local state transitions between configurations of processes, only mentioning the parts of a configuration they rewrite. We use the predicates $\mathsf{proc}(a_m,\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, P_{a_m})$ and $\mathsf{unavail}(a_{\mathsf{S}},\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}})$ to define the states of a configuration (see Sect. 5.1). The former denotes a process executing term P that provides along channel a_m at mode m with worlds $\mathsf{w}_{\mathsf{a}_1}$, $\mathsf{w}_{\mathsf{a}_2}$, and $\mathsf{w}_{\mathsf{a}_3}$ for self, min, and max, respectively. The latter acts as a placeholder for a shared process providing along channel a_{S} with worlds $\mathsf{w}_{\mathsf{a}_1}$, $\mathsf{w}_{\mathsf{a}_2}$, and $\mathsf{w}_{\mathsf{a}_3}$ for self, min, and max, respectively, that is currently unavailable. We note that since worlds are also run-time artifacts, they must occur as part of the state-defining predicates.

Fig. 5 lists selected rules of the dynamics. Since the rules remain largely the same as those of $\mathsf{SILL}_{\mathsf{S}}$, apart from the world annotations that are "threaded through" unchanged, we only discuss the rules that actually differ from the $\mathsf{SILL}_{\mathsf{S}}$ rules. The interested reader can find the remaining rules in the companion technical report [4].

(D-Spawn$_{\mathsf{LL}}$)
$\mathsf{proc}(a_{\mathsf{L}},\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, x_{\mathsf{L}} : A_{\mathsf{L}}[\mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}] \leftarrow X_{\mathsf{L}} \leftarrow \overline{c_{\mathsf{L}}}, \overline{\tilde{c}_{\mathsf{L}}}, \overline{d_{\mathsf{S}}}\,;\ Q_{x_{\mathsf{L}}}),$
$!\mathsf{def}(\Psi' \vdash x'_{\mathsf{L}} : A'_{\mathsf{L}}[\delta_j \updownarrow^{\delta_n}_{\delta_k}] \leftarrow X_{\mathsf{L}} \leftarrow \Delta', \Phi', \Gamma' = P_{x'_{\mathsf{L}}, \mathsf{dom}(\Delta'), \mathsf{dom}(\Phi'), \mathsf{dom}(\Gamma'), \Psi''})$
$\longrightarrow\ \mathsf{proc}(b_{\mathsf{L}},\, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}},\, [b_{\mathsf{L}}/x'_{\mathsf{L}}, \overline{c_{\mathsf{L}}}/\mathsf{dom}(\Delta'), \overline{\tilde{c}_{\mathsf{L}}}/\mathsf{dom}(\Phi'), \overline{d_{\mathsf{S}}}/\mathsf{dom}(\Gamma')]\hat{\gamma}(P_{x'_{\mathsf{L}}, \mathsf{dom}(\Delta'), \mathsf{dom}(\Phi'), \mathsf{dom}(\Gamma'), \Psi''})),$
$\quad \mathsf{proc}(a_{\mathsf{L}},\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, [b_{\mathsf{L}}/x_{\mathsf{L}}]Q_{x_{\mathsf{L}}}),$
$\quad \mathsf{unavail}(b_{\mathsf{S}},\, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}) \quad (b \textit{ fresh})$

(D-New)
$\mathsf{proc}(a,\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, \mathsf{w} \leftarrow \mathsf{new_world};\ Q_{\mathsf{w}}) \longrightarrow \mathsf{proc}(a,\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, Q_{\mathsf{w}}) \quad (\mathsf{w} \textit{ fresh})$

(D-Ord)
$\mathsf{proc}(a,\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, \mathsf{w} < \mathsf{w}';\ Q) \longrightarrow \mathsf{proc}(a,\, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}},\, Q)$

Fig. 5. Selected multiset rewriting rules of $\mathsf{SILL}_{\mathsf{S}^+}$.

Noteworthy are the rules D-New and D-Ord for creating and relating worlds, respectively. Rule D-New creates a fresh world, which will be globally available in the configuration. Rule D-Ord, on the other hand, updates the configuration's order with the pair $\mathsf{w} < \mathsf{w}'$. Rule D-Spawn$_{\mathsf{LL}}$, lastly, substitutes actual worlds for world variables in the body of the spawned process, using the substitution mapping γ defined earlier. It relies on the existence of a corresponding definition predicate for each process definition contained in the signature Σ. We note that the substitution γ in rule D-Spawn$_{\mathsf{LL}}$ instantiates the appropriate world variables in the spawned process P.

4 Extended Example: An Imperative Shared Queue

We now develop a typical imperative-style implementation of a queue that uses a list data structure internally to store the queue's elements and has shared references to the front and the back of the list for concurrent dequeueing and enqueueing, respectively. The session types for the queue and the list are[2]

$$\mathsf{queue}\ A_{\mathsf{S}} = \uparrow_{\mathsf{L}}^{\mathsf{S}} \&\{\mathsf{enq} : \Pi x{:}A_{\mathsf{S}}.\ \downarrow_{\mathsf{L}}^{\mathsf{S}}\mathsf{queue}\ A_{\mathsf{S}},\ \mathsf{deq} : \oplus\{\mathsf{none} : \downarrow_{\mathsf{L}}^{\mathsf{S}}\mathsf{queue}\ A_{\mathsf{S}}, \mathsf{some} : \exists x{:}A_{\mathsf{S}}.\ \downarrow_{\mathsf{L}}^{\mathsf{S}}\mathsf{queue}\ A_{\mathsf{S}}\}\}$$

$$\mathsf{list}\ A_{\mathsf{S}} = \uparrow_{\mathsf{L}}^{\mathsf{S}} \&\{\mathsf{ins} : \Pi x{:}A_{\mathsf{S}}.\ \exists y{:}\mathsf{list}\ A_{\mathsf{S}}.\ \downarrow_{\mathsf{L}}^{\mathsf{S}}\mathsf{list}\ A_{\mathsf{S}},\ \mathsf{del} : \oplus\{\mathsf{none} : \downarrow_{\mathsf{L}}^{\mathsf{S}}\mathsf{list}\ A_{\mathsf{S}}, \mathsf{some} : \exists x{:}A_{\mathsf{S}}.\ \downarrow_{\mathsf{L}}^{\mathsf{S}}\mathsf{list}\ A_{\mathsf{S}}\}$$

The list is implemented in terms of processes *empty* and *elem*, denoting the empty list and a cons cell, respectively. We show the more interesting case of a cons cell (Fig. 6). The queue is defined by processes *head* (Fig. 7) and *queue_proc* (Fig. 8), the latter being the queue's interface to its clients.

$$elem : \{\delta_1 < \delta_2, \delta_2 < \delta_3, \delta_3 < \delta_4 \vdash \mathsf{list}[\delta_1 \updownarrow_{\delta_2}^{\delta_2}] A_{\mathsf{S}}[\delta_3 \updownarrow_{\delta_4}^{\delta_4}] \leftarrow A_{\mathsf{S}}[\delta_3 \updownarrow_{\delta_4}^{\delta_4}], \mathsf{list}[\delta_1 \updownarrow_{\delta_2}^{\delta_2}] A_{\mathsf{S}}[\delta_3 \updownarrow_{\delta_4}^{\delta_4}]\}$$

$$\begin{array}{l}
c[\delta_1 \updownarrow_{\delta_2}^{\delta_2}][\delta_3 \updownarrow_{\delta_4}^{\delta_4}] \leftarrow elem \leftarrow x[\delta_3 \updownarrow_{\delta_4}^{\delta_4}], next[\delta_1 \updownarrow_{\delta_2}^{\delta_2}][\delta_3 \updownarrow_{\delta_4}^{\delta_4}] = \\
\quad c' \leftarrow \mathsf{accept}\ c\ ; \\
\quad \mathsf{case}\ c'\ \mathsf{of} \\
\quad \mid \mathsf{ins} \rightarrow y \leftarrow \mathsf{recv}\ c'\ ;\ n \leftarrow elem \leftarrow y,\ next\ ;\ \mathsf{send}\ c'\ n\ ; \\
\qquad\qquad c \leftarrow \mathsf{detach}\ c'\ ; \\
\qquad\qquad c'' : \mathsf{list}[\delta_1 \updownarrow_{\delta_2}^{\delta_2}] A_{\mathsf{S}}[\delta_3 \updownarrow_{\delta_4}^{\delta_4}] \leftarrow elem \leftarrow x,\ n\ ;\ \mathsf{fwd}\ c\ c'' \\
\quad \mid \mathsf{del} \rightarrow c'.\mathsf{some}\ ;\ \mathsf{send}\ c'\ x\ ; \\
\qquad\qquad c \leftarrow \mathsf{detach}\ c'\ ;\ \mathsf{fwd}\ c\ next
\end{array}$$

Fig. 6. Imperative queue – *elem* process.

We can now define a client (Fig. 8) for the queue, assuming existence of a corresponding shared session type item and a process *item_proc* offering a session of type $\mathsf{item}[\delta_3 \updownarrow_{\delta_4}^{\delta_4}]$. The client instantiates the queue at world δ_b, allowing it to acquire resources at world w_1, which is exactly the world at which process *queue_proc* instantiates the list. Given that the client itself resides at world δ_a, which is smaller than the queue's world δ_b, the client is allowed to acquire the queue, which in turn will acquire the list to satisfy any requests by the client.

The example showcases a paradigmatic use of several collaborators, where collaborators can hold resources while they "talk down" in the tree. In particular, as illustrated in Fig. 9, the clients C_1, C_2, and C_3 compete for resources at world δ_b, i.e., the queue Q. On the other hand, a client C_i collaborates with the queue Q, the list elements L_i, and the items I_i, since they do not overlap in

[2] We adopt polymorphism for the example without formal treatment since it is orthogonal and has been studied for session types in [23,46].

$\mathit{head} : \{\delta_0 < \delta_1, \delta_1 < \delta_2, \delta_2 < \delta_3, \delta_3 < \delta_4 \vdash \mathsf{queue}[\delta_0 \updownarrow^{\delta_1}_{\delta_1}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathsf{list}[\delta_1 \updownarrow^{\delta_2}_{\delta_2}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}],$
$\mathsf{list}[\delta_1 \updownarrow^{\delta_2}_{\delta_2}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}]\}$

$c[\delta_0 \updownarrow^{\delta_1}_{\delta_1}][\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{head} \leftarrow \mathit{front}[\delta_1 \updownarrow^{\delta_2}_{\delta_2}][\delta_3 \updownarrow^{\delta_4}_{\delta_4}],\ \mathit{back}[\delta_1 \updownarrow^{\delta_2}_{\delta_2}][\delta_3 \updownarrow^{\delta_4}_{\delta_4}] =$
$\quad c' \leftarrow \mathsf{accept}\ c\ ;$
$\quad \mathsf{case}\ c'\ \mathsf{of}$
$\quad \mid \mathsf{enq} \rightarrow x \leftarrow \mathsf{recv}\ c'\ ;$
$\qquad\qquad \mathit{back}' \leftarrow \mathsf{acquire}\ \mathit{back}\ ;$
$\qquad\qquad \mathit{back}'.\mathsf{ins}\ ;\ \mathsf{send}\ \mathit{back}'\ x\ ;\ e \leftarrow \mathsf{recv}\ \mathit{back}'\ ;$
$\qquad\qquad \mathit{back} \leftarrow \mathsf{release}\ \mathit{back}'\ ;$
$\qquad\qquad c \leftarrow \mathsf{detach}\ c'\ ;\ c'' : \mathsf{queue}[\delta_0 \updownarrow^{\delta_1}_{\delta_1}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{head} \leftarrow \mathit{front},\ e\ ;\ \mathsf{fwd}\ c\ c''$
$\quad \mid \mathsf{deq} \rightarrow \mathit{front}' \leftarrow \mathsf{acquire}\ \mathit{front}\ ;$
$\qquad\qquad \mathit{front}'.\mathsf{del}\ ;$
$\qquad\qquad (\mathsf{case}\ \mathit{front}'\ \mathsf{of}$
$\qquad\qquad \mid \mathsf{none} \rightarrow \mathit{front} \leftarrow \mathsf{release}\ \mathit{front}'\ ;\ c'.\mathsf{none}\ ;\ c \leftarrow \mathsf{detach}\ c'\ ;$
$\qquad\qquad\qquad c'' : \mathsf{queue}[\delta_0 \updownarrow^{\delta_1}_{\delta_1}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{head} \leftarrow \mathit{front},\ \mathit{back}\ ;\ \mathsf{fwd}\ c\ c''$
$\qquad\qquad \mid \mathsf{some} \rightarrow x \leftarrow \mathsf{recv}\ \mathit{front}'\ ;$
$\qquad\qquad\qquad \mathit{front} \leftarrow \mathsf{release}\ \mathit{front}'\ ;$
$\qquad\qquad\qquad c'.\mathsf{some}\ ;\ \mathsf{send}\ c'\ x\ ;\ c \leftarrow \mathsf{detach}\ c'\ ;$
$\qquad\qquad\qquad c'' : \mathsf{queue}[\delta_0 \updownarrow^{\delta_1}_{\delta_1}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{head} \leftarrow \mathit{front},\ \mathit{back}\ ;\ \mathsf{fwd}\ c\ c'')$

Fig. 7. Imperative queue – *head* process.

$\mathit{queue_proc} : \{\delta_0 < \delta_1, \delta_1 < \delta_3, \delta_3 < \delta_4$
$\quad \vdash \mathsf{queue}[\delta_0 \updownarrow^{\delta_1}_{\delta_1}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}]\}$

$c[\delta_0 \updownarrow^{\delta_1}_{\delta_1}][\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{queue_proc} =$
$\quad \mathsf{w}_2 \leftarrow \mathsf{new_world}\ ;$
$\quad \delta_1 < \mathsf{w}_2\ ;\ \mathsf{w}_2 < \delta_3\ ;$
$\quad e : \mathsf{list}[\delta_1 \updownarrow^{\mathsf{w}_2}_{\mathsf{w}_2}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{empty}\ ;$
$\quad c'' : \mathsf{queue}[\delta_0 \updownarrow^{\delta_1}_{\delta_1}]A_{\mathsf{S}}[\delta_3 \updownarrow^{\delta_4}_{\delta_4}] \leftarrow \mathit{head}$
$\qquad \leftarrow e, e\ ;$
$\quad \mathsf{fwd}\ c\ c''$

$\mathit{client} : \{\delta_a < \delta_b \vdash \mathbf{1}[\delta_a \updownarrow^{\delta_b}_{\delta_b}]\}$

$c[\delta_a \updownarrow^{\delta_b}_{\delta_b}] \leftarrow \mathit{client} =$
$\quad \mathsf{w}_1 \leftarrow \mathsf{new_world}\ ;\ \mathsf{w}_3 \leftarrow \mathsf{new_world}\ ;$
$\quad \mathsf{w}_4 \leftarrow \mathsf{new_world}\ ;$
$\quad \delta_b < \mathsf{w}_1\ ;\ \mathsf{w}_1 < \mathsf{w}_3\ ;\ \mathsf{w}_3 < \mathsf{w}_4\ ;$
$\quad i_0 : \mathsf{item}[\mathsf{w}_3 \updownarrow^{\mathsf{w}_4}_{\mathsf{w}_4}] \leftarrow \mathit{item_proc}\ ;$
$\quad q : \mathsf{queue}[\delta_b \updownarrow^{\mathsf{w}_1}_{\mathsf{w}_1}]A_{\mathsf{S}}[\mathsf{w}_3 \updownarrow^{\mathsf{w}_4}_{\mathsf{w}_4}] \leftarrow \mathit{queue_proc}\ ;$
$\quad q' \leftarrow \mathsf{acquire}\ q\ ;\ q'.\mathsf{enq}\ ;\ \mathsf{send}\ q'\ i_0\ ;$
$\quad q \leftarrow \mathsf{release}\ q'\ ;\ \mathsf{close}\ c$

Fig. 8. Imperative queue – *queue_proc* process and *client* process.

the set of resources they may acquire: a client acquires resources at δ_b, a queue resources at w_1, a list resources at w_2, and an item resources at w_4, and we have $\delta_a < \delta_b < \mathsf{w}_1 < \mathsf{w}_2 < \mathsf{w}_3 < \mathsf{w}_4$. We note in particular that the setup prevents a list element from acquiring its successor, forcing linear access through the queue.

5 Semantics

In this section, we discuss the meta-theoretical properties of $\mathsf{SILL}_{\mathsf{S}^+}$, focusing on deadlock-freedom. The companion technical report [4] provides further details.

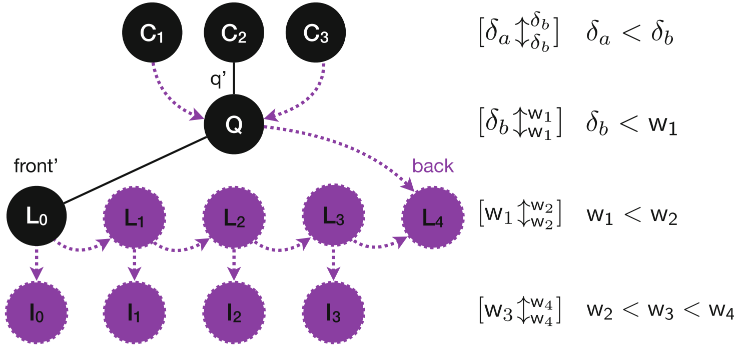

Fig. 9. Run-time process graph for imperative queue (see Fig. 3 for legend).

5.1 Configuration Typing and Preservation

Given the hierarchy between mode S and L and the fact that shared processes cannot depend on linear processes, we divide a configuration into a *shared* part Λ and a linear part Θ. We use the typing judgment $\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta$ to type configurations. The judgment expresses that a well-formed configuration $\Lambda; \Theta$ provides the shared channels in Γ and the linear channels in Φ and Δ. A configuration is type-checked relative to all shared channel references and a global order Ψ. While type-checking is compositional insofar as each process definition can be type-checked separately, solely relying on the process' local Ψ (and Γ), at run-time, the entire order that a configuration relies upon is considered. We give the configuration typing rules in Fig. 10.

Our progress theorem crucially depends on the guarantee that the Invariants 1 and 2 from Sect. 3 hold for every linear process in a configuration's tree. This is expressed by the premises $\mathsf{Inv}_1(\mathsf{proc}(a_\mathsf{L}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_\mathsf{L}}))$ and $\mathsf{Inv}_2(\mathsf{proc}(a_\mathsf{L}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_\mathsf{L}}))$ in rule (T-Θ_2), based on the Definitions 1 and 2 below that restate Invariants 1 and 2 for an entire configuration. We note that Invariant 2 is based on the set of all transitive children (i.e., *descendants*) of a process. We formally define the notion of a descendant inductively over a well-typed linear configuration. The interested reader can find the definition in the companion technical report [4].

Invariant 1 ($\mathsf{min}(\mathsf{parent}) \leq \mathsf{self}(\mathsf{acquired_child}) \leq \mathsf{max}(\mathsf{parent})$)**.** *If* $\Psi; \Gamma \vDash \Theta :: \Phi, \Delta$ *and for any* $\mathsf{proc}(a_\mathsf{L}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_\mathsf{L}}) \in \Theta$ *such that* $\Psi; \Gamma; \Phi_1; \Delta_1 \vdash P_{a_\mathsf{L}} :: (a_\mathsf{L} : A_\mathsf{L}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])$*,* $\mathsf{Inv}_1(\mathsf{proc}(a_\mathsf{L}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_\mathsf{L}}))$ *holds if an only if for every acquired resource* $b_\mathsf{L} : B_\mathsf{L}[\mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}] \in \Phi_1$ *it holds that* $\Psi^* \vdash \mathsf{w}_{\mathsf{a}_2} \leq \mathsf{w}_{\mathsf{b}_1} \leq \mathsf{w}_{\mathsf{a}_3}$*. Moreover, if* $P_{a_\mathsf{L}} = x_\mathsf{L} \leftarrow \mathsf{acquire}\, c_\mathsf{S}\, ; Q_{x_\mathsf{L}}$*, for a* $(c_\mathsf{S} : \uparrow^\mathsf{S}_\mathsf{L} C_\mathsf{L}[\mathsf{w}_{\mathsf{c}_1} \updownarrow^{\mathsf{w}_{\mathsf{c}_3}}_{\mathsf{w}_{\mathsf{c}_2}}]) \in \Gamma$*, then, for every acquired resource* $b_\mathsf{L} : B_\mathsf{L}[\mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}] \in \Phi_1$*, it holds that* $\Psi^+ \vdash \mathsf{w}_{\mathsf{b}_1} < \mathsf{w}_{\mathsf{c}_1}$ *and that* $\Psi^* \vdash \mathsf{w}_{\mathsf{a}_2} \leq \mathsf{w}_{\mathsf{c}_1} \leq \mathsf{w}_{\mathsf{a}_3}$.

$$\frac{}{\Psi; \Gamma \vDash (\cdot) :: (\cdot)} \text{ (T-}\Theta_1\text{)}$$

$$\frac{\begin{array}{c}(a_{\mathsf{S}} : \hat{B}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}]) \in \Gamma \qquad \vdash (A_{\mathsf{L}}, \hat{B})\ \mathsf{sesync} \qquad \Psi \vdash A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}]\ \mathsf{type} \\ \Psi^* \vdash \mathsf{w}_{\mathsf{a}_2} \leq \mathsf{w}_{\mathsf{a}_3} \qquad \mathsf{Inv}_1(\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}})) \qquad \mathsf{Inv}_2(\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}})) \\ \Psi; \Gamma \vDash \Theta :: \Phi, \Phi_1, \Delta, \Delta_1 \qquad \Psi; \Gamma; \Phi_1; \Delta_1 \vdash P_{a_{\mathsf{L}}} :: (a_{\mathsf{L}} : A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])\end{array}}{\Psi; \Gamma \vDash \Theta, \mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}}) :: (\Phi, \Delta, a_{\mathsf{L}} : A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])} \text{ (T-}\Theta_2\text{)}$$

$$\frac{}{\Psi; \Gamma \vDash (\cdot) :: (\cdot)} \text{ (T-}\Lambda_1\text{)} \qquad \frac{\begin{array}{c}\vdash (\uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}, \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}})\ \mathsf{sesync} \qquad \Psi \vdash \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}]\ \mathsf{type} \\ \Psi^* \vdash \mathsf{w}_{\mathsf{a}_2} \leq \mathsf{w}_{\mathsf{a}_3} \qquad \Psi; \Gamma \vdash P_{a_{\mathsf{S}}} :: (a_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])\end{array}}{\Psi; \Gamma \vDash \mathsf{proc}(a_{\mathsf{S}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{S}}}) :: (a_{\mathsf{S}} : \uparrow^{\mathsf{S}}_{\mathsf{L}} A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])} \text{ (T-}\Lambda_2\text{)}$$

$$\frac{}{\Psi; \Gamma \vDash \mathsf{unavail}(a_{\mathsf{S}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}) :: (a_{\mathsf{S}} : \hat{A}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])} \text{ (T-}\Lambda_3\text{)} \qquad \frac{\Psi; \Gamma \vDash \Lambda :: \Gamma_1 \qquad \Psi; \Gamma \vDash \Lambda' :: \Gamma_2}{\Psi; \Gamma \vDash \Lambda, \Lambda' :: \Gamma_1, \Gamma_2} \text{ (T-}\Lambda_4\text{)}$$

$$\frac{\Psi; \Gamma \vDash \Lambda :: \Gamma \qquad \Psi; \Gamma \vDash \Theta :: \Phi, \Delta}{\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta} \text{ (T-}\Omega\text{)}$$

Fig. 10. Configuration typing

Invariant 2 (max(parent) $<$ minima(descendants)). *If* $\Psi; \Gamma \vDash \Theta :: \Phi, \Delta$ *and for any* $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}}) \in \Theta$ *and that process' descendants* $(\Psi; \Gamma \vDash \Theta :: \Phi, \Delta) \rhd a_{\mathsf{L}} = (\Phi', \Delta')$, $\mathsf{Inv}_2(\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}}))$ *holds iff for every descendant* $b_{\mathsf{L}} : B_{\mathsf{L}}[\mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}] \in (\Phi', \Delta')$ *it holds that* $\Psi^+ \vdash \mathsf{w}_{\mathsf{a}_3} < \mathsf{w}_{\mathsf{b}_2}$.

Our preservation theorem states that Invariants 1 and 2 are preserved for every linear process in the configuration along transitions. Moreover, the theorem expresses that the types of the providing linear channels Φ and Δ are maintained along transitions and that new shared channels and worlds may be allocated. The proof relies, in particular, on session types being strictly equi-synchronizing, on a process' type well-formedness and assurance that the process' min world is less than or equal to its max world.

Theorem 5.1 (Preservation). *If* $\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta$ *and* $\Lambda; \Theta \longrightarrow \Lambda'; \Theta'$, *then* $\Psi'; \Gamma' \vDash \Lambda'; \Theta' :: \Gamma'; \Phi, \Delta$, *for some* Λ', Θ', Ψ', *and* Γ'.

5.2 Progress

In our development so far we have distilled the two scenarios of interdependencies between processes that can lead to deadlocks: *cyclic acquisitions* and *interdependent acquisitions and synchronizations*. This has lead to the development of a type system that ingrains the notions of *competitors* and *collaborators*, such that the former compete for a set of resources whereas the latter do not overlap in the set of resources they acquire. Our type system then ties these notions to a configuration's linear process tree such that collaborators stand in a parent-descendant relationship to each other and competitors in a sibling/cousin relationship. In this section, we prove that this orchestration is sufficient to rule out any of the aforementioned interdependencies.

To this end we introduce the notions of *red* and *green arrows* that allow us to reason about process interdependencies in a configuration's tree. A red arrow points from a linear $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, Q)$ to a linear $\mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, P)$, if the former is attempting to acquire a resource held by the latter and, consequently, is waiting for the latter to release that resource. A green arrow points from a linear $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, Q)$ to a linear $\mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, P)$, if the former is waiting to synchronize with the latter. We define these arrows formally as follows:

Definition 5.2 (Acquire Dependency — "Red Arrow"). *Given a well-formed and well-typed configuration $\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta$, there exists a waiting-due-to-acquire relation $\mathcal{A}(\Theta)$ among linear processes in Θ at run-time such that*

$$\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, x_{\mathsf{L}} \leftarrow \mathsf{acquire}\ c_{\mathsf{S}};\ Q_{x_{\mathsf{L}}}) <_{\mathcal{A}} \mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, P\langle c_{\mathsf{L}} \rangle)$$

where $P\langle c_{\mathsf{L}} \rangle$ denotes a process term with an occurrence of channel c_{L}.

Definition 5.3 (Synchronization Dependency — "Green Arrow"). *Given a well-formed and well-typed configuration $\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta$, there exists a waiting-due-to-synchronization relation $\mathcal{S}(\Theta)$ among linear processes in Θ at run-time such that*

$$\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, _\langle b_{\mathsf{L}} \rangle; Q) <_{\mathcal{S}} \mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, _\langle \neg b_{\mathsf{L}} \rangle; P)$$
$$\mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, _\langle b_{\mathsf{L}} \rangle; P) <_{\mathcal{S}} \mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, _\langle \neg b_{\mathsf{L}} \rangle; Q\langle b_{\mathsf{L}} \rangle)$$

where $P\langle a_{\mathsf{L}} \rangle$ denotes a process term with an occurrence of channel b_{L}, $_\langle a \rangle; P$ a process term that currently executes an action along channel a, and $_\langle \neg a \rangle; P$ a process term whose currently executing action does not involve the channel a.

It may be helpful to consult Fig. 3 at this point and note the semantic difference between the violet arrows in that figure and the red arrows discussed here. Whereas violet arrows point from the acquiring process to the resource being acquired, red arrows point from the acquiring process to the process that is holding the resource. Thus, violet arrows can go out of the tree, while red arrows stay within. Given the definitions of red and green arrows, we can define the relation $\mathcal{W}(\Theta)$ on the configuration's tree, which contains all process pairs that are in some way waiting for each other:

Definition 5.4 (Waiting Dependency). *Given a well-formed and well-typed configuration $\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta$, there exists a waiting relation $\mathcal{W}(\Theta)$ among processes in Θ at run-time such that $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P) <_{\mathcal{W}} \mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, Q)$,*

- *if* $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P) <_{\mathcal{A}} \mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, Q)$, *or*
- *if* $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P) <_{\mathcal{S}} \mathsf{proc}(b_{\mathsf{L}}, \mathsf{w}_{\mathsf{b}_1} \updownarrow^{\mathsf{w}_{\mathsf{b}_3}}_{\mathsf{w}_{\mathsf{b}_2}}, Q)$.

Having defined the relation $\mathcal{W}(\Theta)$, we can now state the key lemma underlying our progress theorem, indicating that $\mathcal{W}(\Theta)$ is acyclic in a well-formed and well-typed configuration.

Lemma 5.5 (Acyclicity of $\mathcal{W}(\Theta)$). *If* $\Psi; \Gamma \vDash \Lambda; \Theta :: \Gamma; \Phi, \Delta$*, then* $\mathcal{W}(\Theta)$ *is acyclic.*

We focus on explaining the main idea of the proof here. The proof proceeds by induction on $\Psi; \Gamma \vDash \Theta :: \Phi, \Delta$, assuming for the non-empty case $\Psi; \Gamma \vDash \Theta, \mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}}) :: (\Phi, \Delta, a_{\mathsf{L}} : A_{\mathsf{L}}[\mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}])$ that $\mathcal{W}(\Theta)$ is acyclic, by the inductive hypothesis. We then know that there cannot exist any paths of green and red arrows in Θ that form a cycle, and we have to show that there is no way of introducing such a cyclic path by adding node $\mathsf{proc}(a_{\mathsf{L}}, \mathsf{w}_{\mathsf{a}_1} \updownarrow^{\mathsf{w}_{\mathsf{a}_3}}_{\mathsf{w}_{\mathsf{a}_2}}, P_{a_{\mathsf{L}}})$ to the configuration Θ. In particular, the proof considers all possible new arrows that may be introduced by adding the node and that are necessary for creating a cycle, showing that such arrows cannot come about in a well-typed configuration.

We illustrate the reasoning for the two selected cases shown in Fig. 11. Case **(a)** represents a case in which process $P_{a_{\mathsf{L}}}$ is waiting to synchronize with its child $P_{b_{\mathsf{L}}}$ while holding a resource a descendant of $P_{b_{\mathsf{L}}}$ or $P_{b_{\mathsf{L}}}$ itself wants to acquire. However, this scenario cannot come about in a well-typed configuration because $P_{a_{\mathsf{L}}}$ and $P_{b_{\mathsf{L}}}$ are collaborators and thus cannot overlap in resources they acquire. Case **(b)** represents a case in which process $P_{a_{\mathsf{L}}}$ is waiting to synchronize with its child $P_{b_{\mathsf{L}}}$ while another child, process $P_{c_{\mathsf{L}}}$, is waiting to synchronize with $P_{a_{\mathsf{L}}}$. Given acyclicity of $\mathcal{W}(\Theta)$, a necessary condition for a cycle to form is that there already must exist a red arrow **C** in the configuration that connects the subtrees in which the siblings $P_{b_{\mathsf{L}}}$ and $P_{c_{\mathsf{L}}}$ reside. However, this scenario cannot come about in a well-typed configuration because $P_{b_{\mathsf{L}}}$ and $P_{c_{\mathsf{L}}}$ are competitors, forcing $P_{c_{\mathsf{L}}}$ or any of its descendant to release a resource before synchronizing with $P_{a_{\mathsf{L}}}$. These arguments are made precise in various lemmas in [4].

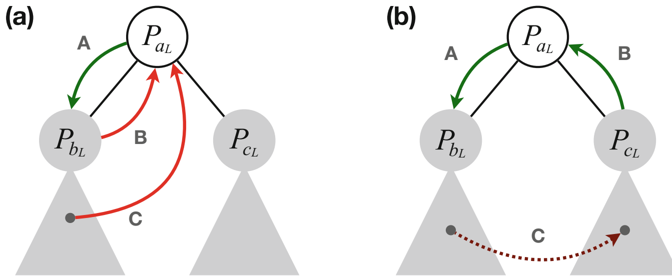

Fig. 11. Two prototypical cases in proof of acyclicty of $\mathcal{W}(\Theta)$.

Given acyclicity of $\mathcal{W}(\Theta)$, we can state and prove the following strong progress theorem. The theorem relies on the notion of a *poised* process, a process currently executing an action along its offering channel, and distinguishes a configuration only consisting of the top-level, linear "main" process from one that consists of several linear processes. We use $|\Theta|$ to denote the cardinality of Θ:

Theorem 5.6 (Progress). *If* $\Psi; \Gamma \vDash \Lambda; \Theta :: (\Gamma; c_{\mathsf{L}} : \mathbf{1}[\mathsf{w}_{\mathsf{c}_1} \updownarrow^{\mathsf{w}_{\mathsf{c}_3}}_{\mathsf{w}_{\mathsf{c}_2}}])$*, then either*

- $\Lambda \longrightarrow \Lambda'$*, for some* Λ'*, or*
- Λ *is poised and*
  - *if* $|\Theta| = 1$*, then either* $\Lambda; \Theta \longrightarrow \Lambda'; \Theta'$*, for some* Λ' *and* Θ'*, or* Θ *is poised, or*
  - *if* $|\Theta| > 1$*, then* $\Lambda; \Theta \longrightarrow \Lambda'; \Theta'$*, for some* Λ' *and* Θ'*.*

The theorem indicates that, as long as there exist at least two linear processes in the configuration, the configuration can always step. If the configuration only consists of the main process, then this process will become poised (i.e., ready to close), once all sub-computations are finished. The proof of the theorem relies on the acyclicity of $\mathcal{W}(\Theta)$ and the fact that all sessions must be strictly equi-synchronizing.

6 Additional Discussion

Linear Forwarding. Our current formalization does not include linear forwarding because a forward changes the process tree and thus endangers the invariants imposed on it. This means that certain programs from the purely linear fragment may not type-check in our system. However, the correspondingly η-expanded versions of these programs should be expressible and type-checkable in $\mathsf{SILL}_{\mathsf{S}^+}$. As part of future work, we want to explore the addition of the linear forward

$$\frac{\Psi^+ \vdash \omega_n < \omega_u}{\Psi;\ \Gamma;\ \cdot;\ y_{\mathsf{L}} : A_{\mathsf{L}}[\omega_m \updownarrow^{\omega_v}_{\omega_u}] \vdash \mathsf{fwd}\ x_{\mathsf{L}}\, y_{\mathsf{L}} :: (x_{\mathsf{L}} : A_{\mathsf{L}}[\omega_j \updownarrow^{\omega_n}_{\omega_k}])}\ (\text{T-Id}_{\mathsf{L}})$$

which allows forwarding to processes that are known to not yet be aliased and whose world annotations meet the premise $\Psi^+ \vdash \omega_n < \omega_u$. Restricting to processes in Δ should uphold Invariant 1, while the premise of the rule should uphold Invariant 2. However, this change will affect the inner working of the proofs, the use of inversion in particular, which might have far-reaching consequences that need to be carefully explored.

Unbounded Process Networks and World Polymorphism. The typing discipline presented in the previous sections, while rich enough to account for a wide range of interesting programs, cannot type programs that spawn a statically undetermined number of shared sessions that are then to be used. For instance, while we can easily type a configuration of any given number of dining philosophers (Sect. 3.3), we cannot type a recursive process in which the number of philosophers (and forks) is potentially unbounded (as done in [21,38]), due to the way worlds are created and propagated across processes.

The general issue lies in implementing a statically unbounded network of processes that interact with each other. These interactions require the processes to be spawned at different worlds which must be generated dynamically as needed.

To interact with such a statically unknown number of processes uniformly, their offering channels must be stored in a list-like structure for later use. However, in our system, recursive types have to be invariant with respect to worlds. For instance, in a recursive type such as $T = A_{\mathsf{L}}@\omega_l \updownarrow_{\omega_p}^{\omega_r} \otimes T$, the worlds ω_l, ω_p, ω_r are fixed in the unfoldings of T. Thus, we cannot type a world-heterogeneous list and cannot form such process networks.

Given that the issues preventing us from typing such unbounded networks lie in problems of world invariance, the natural solution is to explore some form of *world polymorphism*, where types can be parameterized by worlds which are instantiated at a later stage. Such techniques have been studied in the context of hybrid logical processes in [7] by considering session types of the form $\forall\delta.A$ and $\exists\delta.A$, sessions that are parametric in the world variable δ, that is instantiated by a concrete reachable world at runtime. While their development cannot be mapped directly to our setting, it is a promising avenue of future work.

7 Related Work

Behavioral Type Analysis of Deadlocks. The addition of channel usage information to types in a concurrent, message-passing setting was pioneered by Kobayashi and Igarashi [30,34], who applied the idea to deadlock prevention in the π-calculus and later to more general properties [31,32], giving rise to a generic system that can be instantiated to produce a variety of concrete typing disciplines for the π-calculus (e.g., race detection, deadlock detection, etc.).

This line of work types π-calculus processes with a simplified form of *process* (akin to CCS [42] terms without name restriction) that characterizes the input/output behavior of processes. These types are augmented with abstract data that pertain to the relative ordering of channel actions, with the type system ensuring that the transitive closure of such orderings forms a strict partial order, ensuring deadlock-freedom (i.e., communication succeeds unless a process diverges). Building on this, Kobayashi et al. proposed type systems that ensure a stronger property dubbed lock-freedom [35] (i.e., communication always succeeds), and variants that are amenable to type inference [36,39]. Kobayashi [37] extended this latter system to more accurately account for recursive processes while preserving the existence of a type inference algorithm.

Our system draws significant inspiration from this line of work, insofar as we also equip types with abstract ordering data on certain communication actions, which is then statically enforced to form a strict partial order. We note that our $\mathsf{SILL}_{\mathsf{S}^+}$ language differs sufficiently from the pure π-calculus in terms of its constructs and semantics to make the formulation of a direct comparison or an immediate application of their work unclear (e.g., [37] uses replication to encode recursive processes). Moreover, we integrate this style of order-based reasoning with both linear and shared session typing, which interact in non-trivial ways (especially in the presence of recursive types and recursive process definitions).

In terms of typability, enforcing session fidelity can be a double-edged sword: some examples of the works above can be transposed to $\mathsf{SILL}_{\mathsf{S}^+}$ with mostly

cosmetic changes and without making use of shared sessions (e.g., a parallel implementation of factorial that recurses via replication but always answers on a private channel); others are incompatible with linear sessions and require the use of shared sessions via the acquire-release discipline, which entails a more indirect but still arguably faithful modelling of the original π-calculus behavior; some examples, however, cannot be easily adapted to the shared session discipline (e.g., $*c?(x,y).x?(z).y?(z) \mid *c?(x,y).y?(z).x?(z)$ is typable in [37], where $x?(z)$ denotes input on x and $*c?(x,y)$ denotes replicated input) and their transcription, while possible, would be too far removed from the original term to be deemed a faithful representation. Recursive processes are known to produce patterns that can be challenging to analyze using such order-based techniques. The work of [21,38] specializes Kobayashi's system to account for potentially unbounded process networks with non-trivial forms of sharing. Such systems are not typable in our work (see Sect. 6 for additional discussion on this topic).

The work of Padovani [44] develops techniques inspired by [35,37] to develop a typing system for deadlock (and lock) freedom for the linear π-calculus where (linear) channels must be used exactly once. By enforcing this form of linearity, the resulting system uses only one piece of ordering data per channel usage and can easily integrate a form of channel polymorphism that accounts for intricate cyclic interleavings of recursive processes. The combination of manifest sharing and linear session typing does not seem possible without the use of additional ordering data, and the lack of single-use linear channels make the robust channel polymorphism of [44] not feasible in our setting.

Dardha and Gay [15] recently integrated a system of Kobayashi-style orderings in a logical session π-calculus based on classical linear logic, extended with the ability to form *cyclic dependencies* of actions on *linear* session channels (Atkey et al. [1] study similar cycles but do not consider deadlock-freedom), without the need for new process constructs or an acquire-release discipline. Their work considers only a restricted form of replication common in linear logic-based works, not including recursive types nor recursive process definitions. This reduces the complexity of their system, at the cost of expressiveness. We also note that the cycles enabled by their system are produced by processes sharing multiple *linear* names. Since linearity is still enforced, they cannot represent the more general form of cycles that exploit shared channels, as we do.

A comparative study of session typing and Kobayashi-style systems in terms of sharing was developed by Dardha and Pérez [16], showing that such order-based techniques can account for sharing in ways that are out of reach of both classical session typing and pure logic-based session typing. Our system (and that of [15]) aims to combine the heightened power of Kobayashi-style systems with the benefits of session typing, which seems to be better suited as a typing discipline for a high-level programming language [18].

Progress and Session Typing. To address limitations of classical binary session types, Honda et al. [27] introduced *multiparty* session types, where sessions are described by so-called global types that capture the interactions between an arbitrary number of session participants. Under some well-formedness

constraints, global types can be used to ensure that a collection of processes correctly implements the global behavior in a deadlock-free way. However, these global type-based approaches do not ensure deadlock freedom in the presence of higher-order channel passing or interleaved multiparty sessions. Coppo et al. [13] and Bettini et al. [6] develop systems that track usage orders among interleaved multiparty sessions, ruling out cyclic dependencies that can lead to deadlocks. The resulting system is quite intricate, since it combines the full multiparty session theory with the order tracking mechanism, interacts negatively with recursion (essentially disallowing interleaving with recursion) and, by tracking order at the multiparty session-level, ends up rejecting various benign configurations that can be accounted for by our more fine-grained analysis. We also highlight the analyses of Vieira and Vasconcelos [54] and Padovani et al. [45] that are more powerful than the approaches above, at the cost of a more complex analysis based on conversation types [10] (themselves a partial-order based technique).

Static Analysis of Concurrent Programs. Lange et al. [40,41] develop a deadlock detection framework applied to the Go programming language. Their work distills CCS processes from programs which are then checked for deadlocks by a form of symbolic execution [40] and *model-checked* against modal μ-calculus formulae [41] which encode deadlock-freedom of the abstracted process (among other properties of interest). Their abstraction introduces some distance between the original program and the analysed process and so the analysis is sound only for certain restricted program fragments, excluding any combination of recursion and process spawning. Our direct approach does not suffer from this limitation.

de'Liguoro and Padovani [17] develop a typing discipline for deadlock-freedom in a setting where processes exchange messages via unordered mailboxes. Their calculus subsumes the actor model and their analysis combines both so-called mailbox types and specialized dependency graphs to track potential cycles between mailboxes in actor-based systems. The unordered nature of actor-based communication introduces significant differences wrt our work, which crucially exploits the ordering of exchanged messages.

8 Concluding Remarks

In this paper we have developed the concept of manifest deadlock-freedom in the context of the language $\mathsf{SILL}_{\mathsf{S}^+}$, a shared session-typed language, showcasing both the programming methodology and the expressiveness of our framework with a series of examples. Deadlock-freedom of well-typed programs is established by a novel abstraction of so-called green and red arrows to reason about the interdependencies between processes in terms of linear and shared channel references.

In future work, we plan to address some of the limitations of the interactions of deadlock-free shared sessions with recursion, by considering promising notions of world polymorphism and world communication. We also plan to study the problem of world inference and the inclusion of a linear forwarding construct.

References

1. Atkey, R., Lindley, S., Morris, J.G.: Conflation confers concurrency. In: Lindley, S., McBride, C., Trinder, P., Sannella, D. (eds.) A List of Successes That Can Change the World. LNCS, vol. 9600, pp. 32–55. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30936-1_2
2. Balzer, S., Pfenning, F.: Manifest sharing with session types. Proc. ACM Program. Lang. (PACMPL) **1**(ICEP), 37:1–37:29 (2017)
3. Balzer, S., Pfenning, F., Toninho, B.: A universal session type for untyped asynchronous communication. In: 29th International Conference on Concurrency Theory (CONCUR). LIPIcs, pp. 30:1–30:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018)
4. Balzer, S., Toninho, B., Pfenning, F.: Manifest deadlock-freedom for shared session types. Technical report CMU-CS-19-102, Carnegie Mellon University (2019)
5. Benton, P.N.: A mixed linear and non-linear logic: proofs, terms and models. In: Pacholski, L., Tiuryn, J. (eds.) CSL 1994. LNCS, vol. 933, pp. 121–135. Springer, Heidelberg (1995). https://doi.org/10.1007/BFb0022251
6. Bettini, L., Coppo, M., D'Antoni, L., Luca, M.D., Dezani-Ciancaglini, M., Yoshida, N.: Global progress in dynamically interleaved multiparty sessions. In: van Breugel, F., Chechik, M. (eds.) CONCUR 2008. LNCS, vol. 5201, pp. 418–433. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85361-9_33
7. Caires, L., Pérez, J.A., Pfenning, F., Toninho, B.: Logic-based domain-aware session types, unpublished draft
8. Caires, L., Pfenning, F.: Session types as intuitionistic linear propositions. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 222–236. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15375-4_16
9. Caires, L., Pfenning, F., Toninho, B.: Linear logic propositions as session types. Math. Struct. Comput. Sci. **26**(3), 367–423 (2016)
10. Caires, L., Vieira, H.T.: Conversation types. Theor. Comput. Sci. **411**(51–52), 4399–4440 (2010)
11. Castro, D., Hu, R., Jongmans, S., Ng, N., Yoshida, N.: Distributed programming using role-parametric session types in go: statically-typed endpoint APIs for dynamically-instantiated communication structures. PACMPL **3**(POPL), 29:1–29:30 (2019)
12. Cervesato, I., Scedrov, A.: Relating state-based and process-based concurrency through linear logic. Inf. Comput. **207**(10), 1044–1077 (2009)
13. Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L.: Global progress for dynamically interleaved multiparty sessions. Math. Struct. Comput. Sci. **26**(2), 238–302 (2016)
14. Crary, K., Harper, R., Puri, S.: What is a recursive module? In: ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), pp. 50–63 (1999)
15. Dardha, O., Gay, S.J.: A new linear logic for deadlock-free session-typed processes. In: Baier, C., Dal Lago, U. (eds.) FoSSaCS 2018. LNCS, vol. 10803, pp. 91–109. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89366-2_5
16. Dardha, O., Pérez, J.A.: Comparing deadlock-free session typed processes. In: EXPRESS/SOS, pp. 1–15 (2015)
17. de'Liguoro, U., Padovani, L.: Mailbox types for unordered interactions. In: 32nd European Conference on Object-Oriented Programming, ECOOP 2018, pp. 15:1–15:28 (2018)

18. Gay, S.J., Gesbert, N., Ravara, A.: Session types as generic process types. In: 21st International Workshop on Expressiveness in Concurrency and 11th Workshop on Structural Operational Semantics, EXPRESS/SOS 2014, pp. 94–110 (2014)
19. Gay, S.J., Hole, M.: Subtyping for session types in the π-calculus. Acta Informatica **42**(2–3), 191–225 (2005)
20. Gay, S.J., Vasconcelos, V.T., Ravara, A., Gesbert, N., Caldeira, A.Z.: Modular session types for distributed object-oriented programming. In: 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pp. 299–312 (2010)
21. Giachino, E., Kobayashi, N., Laneve, C.: Deadlock analysis of unbounded process networks. In: Baldan, P., Gorla, D. (eds.) CONCUR 2014. LNCS, vol. 8704, pp. 63–77. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44584-6_6
22. Gommerstadt, H., Jia, L., Pfenning, F.: Session-typed concurrent contracts. In: Ahmed, A. (ed.) ESOP 2018. LNCS, vol. 10801, pp. 771–798. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89884-1_27
23. Griffith, D.: Polarized substructural session types. Ph.D. thesis, University of Illinois at Urbana-Champaign (2016)
24. Griffith, D., Pfenning, F.: SILL (2015). https://github.com/ISANobody/sill
25. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 509–523. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-57208-2_35
26. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0053567
27. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pp. 273–284. ACM (2008)
28. Hu, R., Yoshida, N.: Hybrid session verification through endpoint API generation. In: Stevens, P., Wąsowski, A. (eds.) FASE 2016. LNCS, vol. 9633, pp. 401–418. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49665-7_24
29. Hu, R., Yoshida, N.: Explicit connection actions in multiparty session types. In: Huisman, M., Rubin, J. (eds.) FASE 2017. LNCS, vol. 10202, pp. 116–133. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54494-5_7
30. Igarashi, A., Kobayashi, N.: Type-based analysis of communication for concurrent programming languages. In: Van Hentenryck, P. (ed.) SAS 1997. LNCS, vol. 1302, pp. 187–201. Springer, Heidelberg (1997). https://doi.org/10.1007/BFb0032742
31. Igarashi, A., Kobayashi, N.: A generic type system for the Pi-calculus. In: Conference Record of POPL 2001: The 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 128–141 (2001)
32. Igarashi, A., Kobayashi, N.: A generic type system for the Pi-calculus. Theor. Comput. Sci. **311**(1–3), 121–163 (2004)
33. Jespersen, T.B.L., Munksgaard, P., Larsen, K.F.: Session types for rust. In: 11th ACM SIGPLAN Workshop on Generic Programming, WGP 2015, pp. 13–22 (2015)
34. Kobayashi, N.: A partially deadlock-free typed process calculus. In: Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science, pp. 128–139 (1997)
35. Kobayashi, N.: A type system for lock-free processes. Inf. Comput. **177**(2), 122–159 (2002)
36. Kobayashi, N.: Type-based information flow analysis for the π-calculus. Acta Inf. **42**(4–5), 291–347 (2005)

37. Kobayashi, N.: A new type system for deadlock-free processes. In: Baier, C., Hermanns, H. (eds.) CONCUR 2006. LNCS, vol. 4137, pp. 233–247. Springer, Heidelberg (2006). https://doi.org/10.1007/11817949_16
38. Kobayashi, N., Laneve, C.: Deadlock analysis of unbounded process networks. Inf. Comput. **252**, 48–70 (2017)
39. Kobayashi, N., Saito, S., Sumii, E.: An implicitly-typed deadlock-free process calculus. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 489–504. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-44618-4_35
40. Lange, J., Ng, N., Toninho, B., Yoshida, N.: Fencing off go: liveness and safety for channel-based programming. In: 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pp. 748–761. ACM (2017)
41. Lange, J., Ng, N., Toninho, B., Yoshida, N.: A static verification framework for message passing in go using behavioural types. In: Proceedings of the 40th International Conference on Software Engineering, ICSE 2018, Gothenburg, Sweden, 27 May–03 June 2018, pp. 1137–1148 (2018)
42. Milner, R.: A Calculus of Communicating Systems. LNCS, vol. 92. Springer, Heidelberg (1980). https://doi.org/10.1007/3-540-10235-3
43. Neykova, R., Hu, R., Yoshida, N., Abdeljallal, F.: A session type provider: compile-time API generation of distributed protocols with refinements in F#. In: Proceedings of the 27th International Conference on Compiler Construction, CC 2018, pp. 128–138 (2018)
44. Padovani, L.: Deadlock and lock freedom in the linear π-calculus. In: Computer Science Logic - Logic in Computer Science (CSL-LICS), pp. 72:1–72:10 (2014)
45. Padovani, L., Vasconcelos, V.T., Vieira, H.T.: Typing liveness in multiparty communicating systems. In: Kühn, E., Pugliese, R. (eds.) COORDINATION 2014. LNCS, vol. 8459, pp. 147–162. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-43376-8_10
46. Pérez, J.A., Caires, L., Pfenning, F., Toninho, B.: Linear logical relations and observational equivalences for session-based concurrency. Inf. Comput. **239**, 254–302 (2014)
47. Pfenning, F., Griffith, D.: Polarized substructural session types. In: Pitts, A. (ed.) FoSSaCS 2015. LNCS, vol. 9034, pp. 3–22. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46678-0_1
48. Reed, J.: A judgmental deconstruction of modal logic, January 2009. http://www.cs.cmu.edu/~jcreed/papers/jdml.pdf, unpublished manuscript
49. Scalas, A., Dardha, O., Hu, R., Yoshida, N.: A linear decomposition of multiparty sessions for safe distributed programming. In: 31st European Conference on Object-Oriented Programming, ECOOP 2017, pp. 24:1–24:31 (2017)
50. Scalas, A., Yoshida, N.: Lightweight session programming in scala. In: 30th European Conference on Object-Oriented Programming, ECOOP 2016, pp. 21:1–21:28 (2016)
51. Toninho, B.: A logical foundation for session-based concurrent computation. Ph.D. thesis, Carnegie Mellon University and New University of Lisbon (2015)
52. Toninho, B., Caires, L., Pfenning, F.: Higher-order processes, functions, and sessions: a monadic integration. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 350–369. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-37036-6_20

53. Vasconcelos, V.T.: Fundamentals of session types. Inf. Comput. **217**, 52–70 (2012)
54. Vieira, H.T., Vasconcelos, V.T.: Typing progress in communication-centred systems. In: De Nicola, R., Julien, C. (eds.) COORDINATION 2013. LNCS, vol. 7890, pp. 236–250. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38493-6_17
55. Wadler, P.: Propositions as sessions. In: 17th ACM SIGPLAN International Conference on Functional Programming (ICFP), pp. 273–286. ACM (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.