Interconnectability of Session-Based Logical Processes

BERNARDO TONINHO, NOVA-LINCS, Universidade Nova de Lisboa

and Imperial College London, Portugal

NOBUKO YOSHIDA, Imperial College London, United Kingdom

In multiparty session types, interconnection networks identify which roles in a session engage in communication (i.e., two roles are connected if they exchange a message). In session-based interpretations of linear logic the analogue notion corresponds to determining which processes are composed, or cut, using compatible channels typed by linear propositions. In this work, we show that well-formed interactions represented in a session-based interpretation of classical linear logic (CLL) form strictly less-expressive interconnection networks than those of a multiparty session calculus. To achieve this result, we introduce a new compositional synthesis property dubbed partial multiparty compatibility (PMC), enabling us to build a global type denoting the interactions obtained by iterated composition of well-typed CLL threads. We then show that CLL composition induces PMC global types without circular interconnections between three (or more) participants. PMC is then used to define a new CLL composition rule that can form circular interconnections but preserves the deadlock-freedom of CLL.

CCS Concepts: • **Theory of computation** → **Distributed computing models; Process calculi**; *Linear logic*; • **Software and its engineering** → *Message passing; Concurrent programming languages; Concurrent programming structures*;

Additional Key Words and Phrases: Session types, classical linear logic, multiparty sessions, synthesis

ACM Reference format:

Bernardo Toninho and Nobuko Yoshida. 2018. Interconnectability of Session-Based Logical Processes. *ACM Trans. Program. Lang. Syst.* 40, 4, Article 17 (December 2018), 42 pages.

<https://doi.org/10.1145/3242173>

1 INTRODUCTION

The discovery of linear logic [28] and the early studies of its connections with concurrent processes [1, 2, 4] can be seen as the origin of a Curry-Howard correspondence for linear logic with typed interactive behaviours, which have led to the developments connecting linear logic and (binary) session types [8]. The understanding of linear logic propositions as session types [29], proofs as concurrent processes and proof simplification as communication not only has produced new

The authors thank the anonymous reviewers, Alceste Scalas and Julien Lange for their extensive comments and suggestions. This work is supported by NOVA LINCS (UID/CEC/04516/2013), EPSRC EP/K034413/1, EP/K011715/1, EP/L00058X/1, EP/N027833/1 and EP/N028201/1.

Authors' addresses: B. Toninho, Departamento de Informática, Gab. 239, Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Quinta da Torre, 2829-516 Caparica, Portugal; email: bttoninho@fct.unl.pt; N. Yoshida, Department of Computing, Imperial College London, 180 Queen's Gate, Room 556, South Kensington Campus, London SW7 2AZ, United Kingdom; email: n.yoshida@imperial.ac.uk.

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

2018 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 0164-0925/2018/12-ART17 \$15.00

<https://doi.org/10.1145/3242173>

Fig. 1. Interconnection Networks of (1).

logically motivated techniques for reasoning about concurrent processes [51] but is also usable to articulate idioms of interaction with strong communication safety guarantees such as *protocol fidelity* and *deadlock freedom* [58]. The logical foundation of session types has also sparked a renewed interest on the theory and practice of session types [11, 14, 35, 39, 40].

The term “interconnection networks” in this article originates from Abramsky et al. [2], denoting the connections between parallel threads (processes) by means of *linear* ports or channels. In standard linear logic-based session frameworks, processes communicate through a session channel that connects exactly two distinct subsystems typed by dual propositions: When one party sends, the other receives; when one party offers a selection, the other chooses. Sessions may be dynamically exchanged via a session name or created by invocation of replicated servers. A combination of these features enables the modelling of complex behaviours between an arbitrary number of concurrent threads. However, the linear typing discipline induced by linear logic enforces very strong separation properties on the interconnections of processes: Composition identifies a process by a single of its used *linear* channels, requiring all other linear channels in the composed processes to be disjoint and implemented by strictly *separate* processes. It is from this property that deadlock-freedom arises in a simple typing discipline at the cost of disallowing more interesting interconnection networks.

This article provides a fresh look at session-based logical processes, based on concepts originating in *multiparty session types*. Motivated by an industry need [55] to specify protocols with more than two interconnected, interacting parties, the multiparty session types framework [31] develops a methodology where types implicitly describe connections between *many* communicating processes. The key idea of the framework consists of taking a *global type* (i.e., a global description of the multiparty interaction), from which we *generate* (or project) *local types* for each communicating party (specifying its interactions with all others parties) and check that each process adheres to its local type. Once all processes are typechecked, their composition can interact without deadlock, following the given global type. Recent work develops the connections of multiparty session types and communicating automata [6, 22, 37, 38], denotational semantics [21], Petri Nets [24], applications to, e.g., secure information flow analysis [10, 16], dynamic monitoring [5], and reversible computing [15]. Multiparty session types have also been integrated into mainstream programming languages such as MPI [41, 47], Java [32, 33, 42, 56], Python [20, 43, 46], C [49], Go [48], Erlang [25, 45, 57], Scala [54], and F# [44].

In multiparty sessions, interconnection networks identify which roles in a session engage in direct communication. Participant p is connected to another participant q iff p may exchange a message with q (or vice versa). Consider the following three-party interaction specified as a global type G :

$$G = p \rightarrow q:(\text{nat}).p \rightarrow r:(\text{bool}).r \rightarrow q:(\text{str}).\text{end} \quad (1)$$

The type G specifies an interaction where role p sends to roles q and r a natural number and a boolean, respectively, followed by r sending to q a string, inducing the interconnection network depicted in Figure 1(a), realisable in a system where each role is implemented by a separate process.

However, the network of Figure 1(a) is *not* realisable in linear logic-based session frameworks, while those of Figure 1(b) are.

We posit three processes with the behaviour ascribed by G , each implementing one role in the multiparty session:

$$P \vdash pq:A, pr:B \quad Q \vdash pq:A^\perp, qr:C \quad R \vdash pr:B^\perp, qr:C^\perp,$$

where P is the implementation of p with channel pq for communication between p and q and pr for communication between p and r ; Q implements role q using channel pq (dually to P , identified by the use of $A^\perp$) and qr for communication with r , and so on. While each process is individually typable, no three-way composition is typable: Composition in logic-based systems requires that the two processes share a single common channel name, which is then hidden under restriction. When we compose P and Q (hiding channel pq), we obtain a process that shares names pr and qr with R and so cannot be composed with it. We note that such an issue arises regardless of the order in which we choose to compose the processes. In essence, multiparty session types lead to richer connection topologies (e.g., circular connections) than those resulting from the identification of processes with channels during composition, at the cost of requiring global types and projection to ensure deadlock-freedom.

In this work, we make precise the informal argument sketched above by developing a framework based on the theory of multiparty session types (MP) that enables us to reason about connection topologies induced by the session-based interpretation of Classical Linear Logic (CLL).

Our framework is based on the observation that, since multiparty sessions subsume the binary sessions that are primitive in logical formulations of session types, it is possible to interpret typable processes in CLL as MP processes via a *structure-* and *typability-preserving* translation that maps CLL channels to MP channels that are indexed by a role and a destination (i.e., a consistent assignment of action prefixes in CLL to action prefixes in MP, identifying threads or cut-free CLL processes as individual MP session participants). Indeed, our mapping turns out to be canonical up-to bijective renaming.

To reason about the induced connection topologies, we build on the *synthesis* approaches to multiparty sessions [36, 37] that invert the projection-based proposals: Instead of starting from a global type and then producing the certifiably deadlock-free local communication specifications, the works based on synthesis take a collection of local specifications (i.e., types) and study the conditions, dubbed *multiparty compatibility*, under which the local views form a deadlock-free global interaction. Our work extends these approaches by introducing a *compositional*, or partial, notion of multiparty compatibility (PMC), which allows us to synthesise a global type that represents the interactions obtained by iterated composition of CLL processes, providing the necessary tools to precisely study CLL connection topologies.

As argued above, we establish that process composition in CLL induces PMC global types without circular interconnections between three or more session participants (thus excluding the network of Figure 1(a)). This result extends to other linear logic-based calculi (e.g., ILL [8]), since the fundamental structures induced by composition are the same as that of Figure 1(b), making precise the observation that well-formed interactions in linear logic-based calculi form strictly less expressive interconnections between participants than those of MP.

At a logical level, our observation is justified by the fact that allowing richer links or interconnections between proofs (i.e., processes) generally results in a failure of the cut elimination property, which reflects on the resulting processes as a failure of (global) progress. However, it is not the case that all such interconnection topologies result in deadlocked communication. Thus, given that PMC establishes sufficient conditions to ensure deadlock-freedom, even in the presence of circular interconnection topologies, we consider an extension to the CLL calculus in the form of a

composition rule that can result in circular interconnection topologies, but that by being restricted to those of PMC global types, ensures deadlock-freedom (and, consequently, cut elimination) even in the presence of such richer links between proof objects.

Crucially, in contrast with other works on multiparty sessions and logic [11, 14], our PMC-based extension does not require modifications to the syntax of propositions or CLL processes. Also, previous work [11] gives an encoding (introduced in Reference [7]) of their multiparty calculus into (binary) classical logic by using an *additional* orchestrator that centralises control of all interactions. Our canonical structure-preserving mapping requires no such processes and admits *more* typed representatives of given interconnection networks than existing works [11, 14] (see Example 6.9 and Example 4.13) even *without* the extended composition rule.

We note that as a consequence of our translation, we may use CLL to guarantee deadlock-freedom and termination of a class of MP processes with interleaved sessions, which is *not* normally guaranteed by MP typing systems [30] where only deadlock-freedom on a single session is ensured.

Contributions and Outline:

- We introduce a structure- and typability-preserving translation of typed interactions in a session-based interpretation of CLL, restricted to processes without replication and higher-order channel passing, showing that the translation is *unique* insofar as there exists no other typability-preserving encoding (up to bijective renaming) that maps an individual thread to a single participant (Section 3);
- We develop a compositional synthesis property, PMC (Section 4), which we use to show that the interconnectability of CLL is strictly less expressive than that of a single multiparty session in MP (Section 5);
- We systematically extend our results to the more intricate settings of higher-order channel passing (Section 6) and replication (Section 7), showing that neither feature enriches the interconnectability of CLL;
- We use PMC to develop an extension of CLL process composition dubbed multicut (Section 8) that enables richer interconnection topologies while preserving deadlock-freedom without modifying the types or syntax of CLL. We also show that our extended calculus is able to type a range of known examples from MP.

Our work does not assume a deep familiarity with the session-based interpretations of linear logic, multiparty session types, or multiparty compatibility, providing introductions to the session calculi in Section 2 and to global types and multiparty compatibility in Section 4. An extended discussion of related work is given in Section 9. The appendix lists additional proofs and definitions.

2 PROCESSES, TYPES, AND TYPING SYSTEMS

This section introduces the two calculi used in our work: The binary session calculus CLL typed using the session type interpretation of classical linear logic [9, 60]; and the multiparty session calculus MP [18, 31]. In both settings, the notion of a session consists of a (predetermined) sequence of interactions on a given communication channel.

2.1 Classical Linear Logic (CLL) as Binary Session Types

We give a brief summary of the interpretation of classical linear logic as sessions, consisting of a variant of that of Wadler's CP [60], introduced in Caires et al. [9], using two context regions and without explicit contraction or weakening rules, which are admissible judgmental principles.

Syntax. The syntax of CLL processes $(P, Q, \dots)$ is given below. Channels are ranged over by x, y, z, u, v, w , where we typically use x, y, z for *linear* channels and u, v, w for *shared* or replicated channels.

$$\begin{array}{ll}
 P, Q ::= \bar{x}\langle y \rangle.P \mid x(y).P & \text{Send and receive} \\
 \mid x.l;P \mid x.\text{case}\{l_i : P\}_{i \in I} & \text{Selection and branching} \\
 \mid \mathbf{0} \mid (P \mid Q) & \text{Inaction and parallel} \\
 \mid (\nu x)P \mid !u(y).P & \text{Hiding and replication}
 \end{array}$$

We consider a synchronous calculus with *fresh* (or bound) channel input and output—i.e., all sent channels are fresh by construction as in all works on logical session types (e.g., [8, 9, 40, 60]), following the internal mobility π -calculus [53]. The calculus also includes branching, selection and replication constructs, with the latter allowing us to represent servers as replicated input-guarded processes, where the corresponding matching output processes act as their clients. We write $fn(P)/bn(P)$ for the free/bound channels of P : In $\bar{x}\langle y \rangle.P$ and $(\nu y)P$, y is a binding occurrence. We write $bv(P)$ for the bound variables of P , noting that in $x(y).P$ and $!x(y).P$, y is bound in P . We often omit $\mathbf{0}$.

Below we define the structural congruence for CLL, which is used in the typing system and the reduction semantics.

Definition 2.1 (Structural Congruence for CLL). Structural congruence of CLL processes, written $P \equiv Q$, is the least congruence defined by the following rules:

$$\begin{array}{l}
 P \equiv_{\alpha} Q \Rightarrow P \equiv Q \quad P \mid \mathbf{0} \equiv P \quad P \mid Q \equiv Q \mid P \quad (P \mid Q) \mid R \equiv P \mid (Q \mid R) \\
 (\nu x)(P \mid Q) \equiv (\nu x)P \mid Q \quad x \notin fn(Q) \quad (\nu x)(!x(y).P) \equiv \mathbf{0}.
 \end{array}$$

Reduction. The reduction semantics for CLL, written $P \rightarrow Q$ and defined up to structural congruence $\equiv$, is given below:

$$\begin{array}{ll}
 \bar{x}\langle y \rangle.P \mid x(y).Q & \rightarrow (\nu y)(P \mid Q) \\
 x.l_j;P \mid x.\text{case}\{l_i : Q_i\}_{i \in I} & \rightarrow P \mid Q_j \quad (j \in I) \\
 \bar{x}\langle y \rangle.P \mid !x(y).Q & \rightarrow (\nu y)(P \mid Q) \mid !x(y).Q \\
 P \rightarrow P' & \Rightarrow P \mid Q \rightarrow P' \mid Q \\
 P \rightarrow P' & \Rightarrow (\nu x)P \rightarrow (\nu x)P' \\
 P \equiv P' \wedge P' \rightarrow Q' \wedge Q' \equiv Q & \Rightarrow P \rightarrow Q.
 \end{array}$$

Definition 2.2 (Live Process). A process P is live, written $live(P)$ iff $P \equiv (\nu \tilde{x})(\pi.Q \mid R)$ or $P \equiv (\nu \tilde{x})(\pi; Q \mid R)$ for some R , sequences of names $\tilde{x}$ and a *non-replicated* guarded process $\pi.Q$ or $\pi; Q$, where π is any non-replicated process prefix.

Note how the definition of live process excludes a process of the form $!u(y).P$, which corresponds to a replicated server that has no remaining users.

Types. The syntax of (logical) binary session types A, B is

$$A, B ::= A \otimes B \mid A \wp B \mid \mathbf{1} \mid \perp \mid \oplus\{l_i : A_i\}_{i \in I} \mid \&\{l_i : A_i\}_{i \in I} \mid ?A \mid !A.$$

Following [9, 60], $\otimes$ corresponds to output of a session of type A followed by behaviour B ; $\wp$ to input of A followed by B ; $\oplus$ and $\&$ to selection and branching; $!A$ to replicated channels of type A (i.e., persistent servers) and $?A$ to clients of such servers. The *dual* of A , written $A^\perp$, is defined as (we omit the involutive cases):

$$1^\perp \triangleq \perp, (A \otimes B)^\perp \triangleq A^\perp \wp B^\perp, (\oplus\{l_i : A_i\}_{i \in I})^\perp \triangleq \&\{l_i : A_i^\perp\}_{i \in I}, (!A)^\perp \triangleq ?A^\perp.$$

$$\begin{array}{c}
(\otimes) \frac{P_1 \vdash_{\text{CL}} \Xi; \Delta, y:A \quad P_2 \vdash_{\text{CL}} \Xi; \Delta', x:B}{\bar{x}(y).(P_1 \mid P_2) \vdash_{\text{CL}} \Xi; \Delta, \Delta', x:A \otimes B} \quad (\wp) \frac{P \vdash_{\text{CL}} \Xi; \Delta, y:A, x:B}{x(y).P \vdash_{\text{CL}} \Xi; \Delta, x:A \wp B} \\
\\
(1) \frac{}{\mathbf{0} \vdash_{\text{CL}} \Xi; x:1} \quad (\perp) \frac{P \vdash_{\text{CL}} \Xi; \Delta}{P \vdash_{\text{CL}} \Xi; \Delta, x:\perp} \\
\\
(\oplus) \frac{P \vdash_{\text{CL}} \Xi; \Delta, x:A_j \quad j \in I}{x.l_j.P \vdash_{\text{CL}} \Xi; \Delta, x:\oplus \{l_i : A_i\}_{i \in I}} \quad (\&) \frac{P_1 \vdash_{\text{CL}} \Xi; \Delta, x:A_1 \quad \dots \quad P_n \vdash_{\text{CL}} \Xi; \Delta, x:A_n}{x.\text{case}\{l_i : P_i\}_{i \in I} \vdash_{\text{CL}} \Xi; \Delta, x:\& \{l_i : A_i\}_{i \in I}} \\
\\
(\text{cut}) \frac{P \vdash_{\text{CL}} \Xi; \Delta, x:A \quad Q \vdash_{\text{CL}} \Xi; \Delta', x:A^\perp}{(\nu x)(P \mid Q) \vdash_{\text{CL}} \Xi; \Delta, \Delta'} \\
\\
(!) \frac{P \vdash_{\text{CL}} \Xi; y:A}{!u(y).P \vdash_{\text{CL}} \Xi; u:!A} \quad (?) \frac{P \vdash_{\text{CL}} \Xi, u:A; \Delta}{P\{x/u\} \vdash_{\text{CL}} \Xi; \Delta, x:?A} \\
\\
(\text{copy}) \frac{P \vdash_{\text{CL}} \Xi, u:A; \Delta, x:A}{\bar{u}(x).P \vdash_{\text{CL}} \Xi, u:A; \Delta} \quad (\text{cut}^!) \frac{P \vdash_{\text{CL}} \Xi; x:A \quad Q \vdash_{\text{CL}} \Xi, u:A^\perp; \Delta}{(\nu u)(!u(x).P \mid Q) \vdash_{\text{CL}} \Xi; \Delta}
\end{array}$$

Fig. 2. CLL typing rules.

Typing System. We define the typing system CLL in Figure 2, assigning the usage of channels in P processes to types A, B . The typing judgement is written $P \vdash_{\text{CL}} \Xi; \Delta$, defined *up to structural congruence* $\equiv$ (i.e., we implicitly have that if $P \vdash_{\text{CL}} \Xi; \Delta$ and $P \equiv Q$ then $Q \vdash_{\text{CL}} \Xi; \Delta$), where Δ is a set of hypotheses of the form $x:A$ (not subject to weakening or contraction), where x stands for a free session channel in P and A is a binary session type; and Ξ is a set of hypotheses of the form $u:A$, subject to weakening and contraction, standing for the channels used in P in an unrestricted (or shared) manner. The typing judgement states that process P uses channels according to the session discipline ascribed by Δ and Ξ . We assume all channel names in Δ and Ξ are distinct. We write $\cdot$ for the empty typing environment and Δ, Δ' for the union of Δ and Δ' , only defined when channels in Δ and Δ' are distinct.

Rule $(\otimes)$ accounts for the session output behaviour, typing a channel x with $A \otimes B$ if the process outputs along x a name y that is used in P_1 with behaviour A and, disjointly, P_2 uses x according to B (this strict separation is crucial for deadlock-freedom); dually, rule $(\wp)$ types a channel x with $A \wp B$ if the process performs an input on x of a session channel y such that y is used in the continuation as A , and x as B ; rule (1) types the inactive process with an arbitrary session channel assigned type 1 ; rule $(\perp)$ types the dual behaviour, which just discards the no longer used name; rule $(\oplus)$ types channel x with $\oplus \{l_i:A_i\}_{i \in I}$ by having the process emit a label l_j with $j \in I$, and then using the channel x according to the type A_j in the corresponding branch; dually, rule $(\&)$ types processes that wait for a choice on channel x , with type $\& \{l_i:A_i\}_{i \in I}$, if the process can account for all of the possible choice labels and corresponding behaviours in the type. Thus, the case construct must contain one process P_i using x according to behaviour A_i for each label in the type. Note the additive nature of the rule, where the context Δ is the same in all premises. This enforces that all possible alternative behaviours make use of the *same* session behaviours.

Rule (cut) composes in parallel two processes P and Q , that use channel x with dual types A and $A^\perp$, by hiding x in the composed process in the conclusion of the rule (since no other process may use x). We note that Δ and Δ' are disjoint, so the only common channel between P and Q is x .

The remaining rules define the typing for the replication constructs. Rule (!) types a replicated input channel u as $!A$ if the continuation P uses the input as A without using any other linear channels, which ensures that the replicas of P do not invalidate the linear typing discipline. Rule (?) moves a session channel of type $?A$ to the appropriate shared context Ξ as A , renaming it to u . Rule (copy) types a usage of a shared channel u by sending a fresh channel x along u , which is then used (linearly) as A in the continuation P . Rule (cut¹) allows for composition of shared sessions, provided the processes use no linear channels. Such a rule is needed to ensure that cut elimination holds structurally (i.e., a linear cut between a session channel of type $!A$ and its dual $?A^\perp$ reduces to a cut¹, which will eventually reduce back to a cut). We note that at the level of typable processes, cut¹ can be represented as a cut between rules ! and ?.

PROPOSITION 2.3 (DEADLOCK-FREEDOM IN CLL [9, 60]). *Suppose $P \vdash_{\text{CL}} \cdot; \Delta$, with $\text{live}(P)$ where Δ is either empty or only contains 1 or $\perp$. We have that $P \rightarrow P'$.*

2.2 Multiparty Session (MP) Calculus

Syntax. We introduce the MP calculus of multiparty sessions, where processes P, Q use channels indexed by roles of the multiparty sessions in which they are used. The syntax of processes and channels is given below:

$P, Q ::= c[p](c'); P \mid c[p](x); P$	Send and receive
$\mid c[p] \oplus l; P \mid c[p] \& \{l_i; P_i\}_{i \in I}$	Selection and branching
$\mid \mathbf{0} \mid (P \mid Q) \mid (\nu s)P$	Inaction, parallel, hiding
$c ::= x \mid s[p]$	variable, role-indexed channel.

Role names are identified by p, q, r ; channels are ranged over by s, t ; c denotes channels with role $s[p]$ or variables x .

Local Types. Role indexed channels $s[p]$ in MP are assigned *local types*, ranged over by S, T , denoting the behaviour of each role per channel. Local types are defined as follows:

$$S, T ::= p \uparrow(T); S \mid p \downarrow(T); S \mid \oplus \{l_i; T_i\}_{i \in I} \mid \& \{l_i; T_i\}_{i \in I} \mid \text{end}.$$

The local types $p \uparrow(T); S$ and $p \downarrow(T); S$, which type the send and receive constructs above, denote output to and input from role p of a session channel of type T , followed by behaviour S , respectively. Types $\oplus \{l_i; T_i\}_{i \in I}$ and $\& \{l_i; T_i\}_{i \in I}$, which type the selection and branching constructs, denote the emission (respectively, reception) of a label l_i to (respectively, from) role p , followed by behaviour T_i . Type end denotes no further behaviour. We define the set of roles of local type T , denoted by $\text{roles}(T)$, as the set of all roles occurring in type T .

Partial Projection and Coherence. To define the typing system for MP, we introduce partial projection and coherence. Partial projection takes a local type (that specifies all interactions for a given role) and a role to produce the binary session type [29] that corresponds to the interactions between the role whose behaviour is denoted by the local type and the given role, from the point of view of the former (e.g., if p is the role behaving according to T_p , the projection of T_p for q produces a binary session type describing the interactions between p and q from the perspective of p). Binary session types S, T are given by (by abuse of notation we re-use the same symbols S, T as for local types):

$$\uparrow(T); S \quad \downarrow(T); S \quad \oplus \{l_i; T_i\}_{i \in I} \quad \& \{l_i; T_i\}_{i \in I} \quad \text{end}$$

and their notion of duality $\overline{T}$ is given by $\overline{\uparrow(T); S} \triangleq \downarrow(T); \overline{S}$, $\overline{\downarrow(T); S} \triangleq \uparrow(T); \overline{S}$, $\overline{\oplus \{l_i; T_i\}_{i \in I}} \triangleq \& \{l_i; \overline{T_i}\}_{i \in I}$, $\overline{\& \{l_i; T_i\}_{i \in I}} \triangleq \oplus \{l_i; \overline{T_i}\}_{i \in I}$, $\overline{\text{end}} \triangleq \text{end}$; and the involutive rules for $\downarrow$ and $\&$.

$$\begin{array}{c}
\text{(end)} \frac{\Gamma \text{ end only}}{\mathbf{0} \vdash_{\text{MP}} \Gamma} \quad \text{(send)} \frac{P \vdash_{\text{MP}} \Gamma, c:S}{c[q]\langle c' \rangle; P \vdash_{\text{MP}} \Gamma, c:q\uparrow(T); S, c':T} \quad \text{(recv)} \frac{P \vdash_{\text{MP}} \Gamma, c:S, x:T}{c[q](x); P \vdash_{\text{MP}} \Gamma, c:q\downarrow(T); S} \\
\text{(sel)} \frac{P \vdash_{\text{MP}} \Gamma, c:T_j \quad j \in I}{c[q] \oplus l_j; P \vdash_{\text{MP}} \Gamma, c: \oplus \{l_i : T_i\}_{i \in I}} \quad \text{(bra)} \frac{P_1 \vdash_{\text{MP}} \Gamma, c:T_1 \quad \dots \quad P_n \vdash_{\text{MP}} \Gamma, c:T_n}{c[q] \& \{l_i:P_i\}_{i \in I} \vdash_{\text{MP}} \Gamma, c: \& q\{l_i : T_i\}_{i \in I}} \\
\text{(comp)} \frac{P \vdash_{\text{MP}} \Gamma \quad Q \vdash_{\text{MP}} \Gamma'}{P \mid Q \vdash_{\text{MP}} \Gamma, \Gamma'} \quad \text{(close)} \frac{P \vdash_{\text{MP}} \Gamma, s[p_1] : T_1, \dots, s[p_n] : T_n \quad \text{co}(s[p_1]:T_1, \dots, s[p_n]:T_n)}{(\nu s)P \vdash_{\text{MP}} \Gamma}
\end{array}$$

Fig. 3. MP typing rules.

Definition 2.4 (Partial Projection). Given a local type T , we define its partial projection onto a participant p , written $T \upharpoonright p$, by induction on the structure of T by

$$\begin{array}{ll}
(r\uparrow(S); T) \upharpoonright p = \begin{cases} \uparrow(S); (T \upharpoonright p) & \text{if } p = r \\ T \upharpoonright p & \text{otherwise} \end{cases} & (\oplus r\{l_i:T_i\}_{i \in I}) \upharpoonright p = \begin{cases} \oplus \{l_i:(T_i \upharpoonright p)\}_{i \in I} & \text{if } p = r \\ \sqcup_{i \in I} (T_i \upharpoonright p) & \text{otherwise} \end{cases} \\
(r\downarrow(S); T) \upharpoonright p = \begin{cases} \downarrow(S); (T \upharpoonright p) & \text{if } p = r \\ T \upharpoonright p & \text{otherwise} \end{cases} & (\& r\{l_i:T_i\}_{i \in I}) \upharpoonright p = \begin{cases} \& \{l_i:(T_i \upharpoonright p)\}_{i \in I} & \text{if } p = r \\ \sqcup_{i \in I} (T_i \upharpoonright p) & \text{otherwise,} \end{cases} \\
\text{end} \upharpoonright p = \text{end}
\end{array}$$

where the *merge* $T \sqcup T'$ of T and T' is defined by $T \sqcup T' \triangleq T$; and with $T = \oplus \{l_i : T_i\}_{i \in I}$ and $T' = \oplus \{l'_j : T'_j\}_{j \in J}$,

$$T \sqcup T' \triangleq \oplus (\{l_h : T_h\}_{h \in I \setminus J} \cup \{l'_h : T'_h\}_{h \in J \setminus I} \cup \{l_h : T_h \sqcup T'_h\}_{h \in I \cap J})$$

if $l_h = l'_h$ for each $h \in I \cap J$; and homomorphic for other types (i.e., $\mathcal{T}[T_1] \sqcup \mathcal{T}[T_2] = \mathcal{T}[T_1 \sqcup T_2]$, where $\mathcal{T}$ is a context of local types). $T \sqcup T'$ is undefined otherwise. Partial projection is undefined if merging is undefined.

Merging is needed for two purposes: (1) to check global types well-formedness (i.e., if merge is undefined then the global type is not well formed) and (2) to allow for more typable protocols. Examples of merging can be found in Section 4.1. Coherence ensures that the local types of interacting roles contain the necessary compatible actions (e.g., if the local type for p specifies an emission to q , the local type for q specifies a reception from p [18, 31]) and all the necessary roles are ascribed a type in the context. To define coherence, we introduce session subtyping. We note that the subtyping relation is inverted w.r.t. the “process-oriented” subtyping [26], because, for convenience, we adopt the “channel-oriented” ordering [17]; an analysis of the two subtyping relations is given in Gay [27].

Definition 2.5 (Session Subtyping). We define the subtyping relation between binary session types, $T \leq S$, as the least relation given by the following rules:

$$\begin{array}{c}
\frac{}{\text{end} \leq \text{end}} \quad \frac{\forall i \in I \quad T_i \leq T'_i}{\oplus \{l_i : T_i\}_{i \in I} \leq \oplus \{l_i : T'_i\}_{i \in I \cup J}} \quad \frac{\forall i \in I \quad T_i \leq T'_i}{\& \{l_i : T_i\}_{i \in I \cup J} \leq \& \{l_i : T'_i\}_{i \in I}} \\
\frac{T \leq T' \quad S \leq S'}{\uparrow(T); S \leq \uparrow(T'); S'} \quad \frac{T' \leq T \quad S \leq S'}{\downarrow(T); S \leq \downarrow(T'); S'}.
\end{array}$$

Definition 2.6 (Coherence). Γ is coherent (denoted by $\text{co}(\Gamma)$) iff $s[p]:T_1 \in \Gamma$ and $s[q]:T_2 \in \Gamma$ with $p \neq q$ imply that $T_1 \upharpoonright q \leq \overline{T_2} \upharpoonright p$; and for all $s[p]:T \in \Gamma$ and $q \in \text{roles}(T)$, $\mathcal{A}[q]:T' \in \Gamma$, for some T' .

Typing Rules. We define the typing system MP in Figure 3, assigning the usage of role-indexed channels to local types. The judgement $P \vdash_{\text{MP}} \Gamma$, where Γ is a set of hypotheses of the form $c:T$, denotes that P uses its channels according to Γ . We assume the same notations and conditions for

Γ and Γ, Γ' as for CLL, where $\cdot$ denotes the empty context and Γ, Γ' denotes the disjoint union of Γ and Γ' , defined only when their domains are disjoint.

Rule (end) types the inactive process in a session context containing only terminated sessions (i.e., the context Γ is a collection of assumptions of the form $c_i:\text{end}$). Rule (send) types the emission of a channel endpoint of type T to role q , assigning c the local type $q\uparrow(T); S$, provided the continuation P uses c according to type S . Dually, rule (recv) types the reception of a value of type T , bound to x in the continuation P , sent by q , with type $q\downarrow(T); S$, provided P uses c according to S . Rules (sel) and (bra) are the MP counterparts of rules ($\oplus$) and ($\&$) from CLL (Figure 2), respectively, with the former typing the emission of a label to q and the latter typing the reception of a label from q . Rule (comp) types parallel composition of processes with disjoint session contexts Γ and Γ' . Rule (close) types a *coherent* multiparty session s by hiding the session channel, provided that process P uses $s[p_1]:T_1, \dots, s[p_n]:T_n$ and the corresponding role indices and local types form a coherent typing context.

Reduction. The reduction semantics for MP processes is given below (omitting closure under structural congruence). They are fundamentally identical to the reduction rules of CLL but require not just the session channel to match but also the role assignment to be consistent:

$$\begin{aligned} s[p][q]\langle s'[p'] \rangle; P \mid s[q][p](x); Q &\rightarrow P \mid Q\{s'[p']/x\} \\ s[p][q] \oplus l_j; P \mid s[q][p] \& \{l_i:Q_i\}_{i \in I} &\rightarrow P \mid Q_j \quad (j \in I) \\ P \rightarrow P' &\Rightarrow P \mid Q \rightarrow P' \mid Q \\ P \rightarrow P' &\Rightarrow (\nu s)P \rightarrow P'. \end{aligned}$$

We highlight that, in contrast to CLL, the typing system MP alone does *not* ensure deadlock-freedom, where deadlock-freedom means that all communication actions always eventually fire for processes typed in an empty context. We assume basic value passing, noting that value passing can be encoded with terminated sessions and that henceforth it will be used freely in the rest of the article.

PROPOSITION 2.7 (DEADLOCK IN MP). *There exists a deadlocked process P that is typable in MP, i.e., $P \vdash_{\text{MP}} \emptyset$ does not imply that P is deadlock-free.*

PROOF. Take $P = s[p][r](x); s[p][q]\langle 7 \rangle, Q = s[q][p](x); s[q][r]\langle \text{tt} \rangle$ and $R = s[r][q](x); s[r][p]\langle \text{a} \rangle$. $(\nu s)(P \mid Q \mid R) \vdash_{\text{MP}} \emptyset$, but $P \mid Q \mid R$ is deadlocked. $\square$

3 RELATING THE CLL AND MP SYSTEMS

In this section, we develop one of our main contributions: the connection between the CLL and MP systems. For presentation purposes, we first consider a restriction of CLL *without* name passing and replication, which are addressed in Section 6 and Section 7, respectively. In the following sections, we tacitly make use of value passing, which can be included straightforwardly in the systems of Section 2. To explicate our approach, consider the following CLL typable processes:

$$\begin{aligned} P &\triangleq x\langle 7 \rangle. y(z). x\langle \text{“hello”} \rangle. \mathbf{0} \vdash_{\text{CL}} x:\text{nat} \otimes \text{str} \otimes \mathbf{1}, y:\text{nat} \wp \perp \\ P' &\triangleq y(z). x\langle 7 \rangle. x\langle \text{“hello”} \rangle. \mathbf{0} \vdash_{\text{CL}} x:\text{nat} \otimes \text{str} \otimes \mathbf{1}, y:\text{nat} \wp \perp. \end{aligned}$$

Both P and P' are typable in the same context; however, P first outputs on x , then inputs on y and then outputs on x again, whereas P' flips the order of the first two actions. By the nature of process composition in CLL, both processes can be safely composed with any typable $R_1 \vdash_{\text{CL}} x:\text{nat} \wp \text{str} \wp \perp$ and $R_2 \vdash_{\text{CL}} y:\text{nat} \otimes \mathbf{1}$. We also observe that, since both P and P' are typable in the same context, CLL typing cannot capture *cross-channel sequential dependencies* (i.e., it cannot distinguish orderings of actions on different channels).

We now consider a mapping from CLL to MP. The following processes Q and Q' are hypothetical translations of P and P' . The notation $s[p][q]$ represents a channel in session s with role p and destination q :

$$Q \triangleq s[q][p]\langle 7 \rangle; s[q][r](z); s[q][p]\langle \text{"hello"} \rangle \quad Q' \triangleq s[q][r](z); s[q][p]\langle 7 \rangle; s[q][p]\langle \text{"hello"} \rangle.$$

The processes P and Q above are similar, insofar as both send 7 to a destination (respectively, x and role p), followed by an input (respectively, on y and from r), followed by an output of "hello" to the initial destination. A similar argument can be made for P' and Q' . Despite P and P' having the same types, we have

$$Q \vdash_{\text{MP}} s[q]:p \uparrow(\text{nat}); r \downarrow(\text{nat}); p \uparrow(\text{str}); \text{end} \quad Q' \vdash_{\text{MP}} s[q]:r \downarrow(\text{nat}); p \uparrow(\text{nat}); p \uparrow(\text{str}); \text{end}.$$

By refining channels with role annotations, MP distinguishes orderings of actions on different session *sub-channels* (i.e., the communication links between the several role pairs). More precisely, by grouping the actions of role q along its two session sub-channels $s[q][p]$ and $s[q][r]$ at the type level, we can precisely track the ordering and causal dependencies on the message exchanges between q and p and those with r . Thus, our goal is to find a precise way to systematically map process P to process Q (and P' to Q') and also generate the corresponding local typing in a typability-preserving way.

To relate CLL with MP processes and preserve typability, we proceed as follows:

Mapping 1: $\boxed{P \vdash_{\text{CL}}^\sigma \Delta}$ We define a mapping σ from *session channels* in CLL to *channels indexed with role and destination* in a single MP session, such that given a single-threaded process in CLL (i.e., a cut-free process), we map its channels to role and destination-annotated channels in MP forming a single multiparty session, capturing the cross-channel causal dependencies that are not codified at the level of CLL types.

Mapping 2: $\boxed{\llbracket P \rrbracket_\sigma}$ We generate local type T from a single thread P w.r.t. σ such that $P \vdash_{\text{CL}}^\sigma \Delta$ so that we can translate P in CLL to $\sigma(P)$ typable in MP.

Mapping 3: $\boxed{P \Vdash_\rho^\sigma \Delta; \Gamma}$ We translate the cut (i.e., parallel composition) between two processes in CLL into MP generating a mapping from channels to *session types* Γ with renaming of free and bound names (σ and ρ). This automatically provides a type- and thread-preserving translation $\rho(\sigma(P))$ into MP, which is *unique* up to bijective renaming.

Mapping 1: Preservation of Threads and Typability. Definition 3.1 provides the mapping from session channels in CLL to those in MP. For now, we consider only CLL processes without replication (i.e., typed without uses of rules $!$, $?$, copy and cut¹ – and thus omit Ξ from the CLL typing judgment) and where $A \otimes B$ and $A \wp B$ are restricted to $1 \otimes B$ and $\perp \wp B$, respectively (i.e., no higher-order channel passing, where $1 \otimes B$ can be seen as sending an abstract value of ground type). Moreover, we assume that uses of rule $(\otimes)$ are such that $P_1 \equiv 0$. We lift the restriction on higher-order channel passing in Section 6.

Definition 3.1 (Channel to Role-Indexed Channel Mapping). Let $P \vdash_{\text{CL}} \Delta$ such that the typing derivation does not use the cut rule. We define a channel to (role-)indexed channel mapping σ such that for all $x, y \in \text{fn}(P)$, if $x \neq y$, then $\sigma(x) = s[p][q]$ and $\sigma(y) = s[p][q']$, for some q, q' such that $q \neq q'$, and *unique* s and p (i.e., s and p are the same MP session channel and principal role across the entire mapping σ). We reject reflexive role assignments of the form $s[p][p]$.

We write $P \vdash_{\text{CL}}^\sigma \Delta$ to denote such a mapping and $c_\sigma(x)$, $p_\sigma(x)$, and $d_\sigma(x)$ to denote the channel, first (principal) and second (destination) roles in the image of x in σ .

$$\begin{array}{c}
\text{(thread)} \frac{P \vdash_{\text{CL}}^{\sigma} \Delta}{P \Vdash_{\emptyset}^{\sigma} \Delta; s[p_{\sigma}]:\llbracket P \rrbracket_{\sigma}} \quad \text{(comp)} \frac{P \vdash_{\text{CL}}^{\sigma} \Delta, x:A \quad Q \Vdash_{\rho}^{\sigma'} \Delta', x:A^{\perp}; \Gamma \quad (\star)}{(\nu x)(P \mid Q) \Vdash_{\rho'}^{(\sigma' \cup \sigma) \setminus \{x\}} \Delta, \Delta'; \Gamma, s[p_{\sigma}]:\llbracket P \rrbracket_{\sigma}} \\
\hline
(\star) \quad \begin{array}{l}
\text{(a) bound channels: } \rho' = \rho \cup (x, s[p_{\sigma}][d_{\sigma}(x)]) \\
\text{(b) role/destination match: } p_{\sigma}(x) = d_{\sigma'}(x) \wedge d_{\sigma}(x) = p_{\sigma'}(x) \\
\text{(c) unique destinations: } \forall z \in \Delta, y \in \Delta'. d_{\sigma}(z) \neq d_{\sigma'}(y) \wedge d_{\sigma}(z), d_{\sigma'}(y) \notin \rho
\end{array}
\end{array}$$

Fig. 4. Parallel composition mapping.

A mapping according to Definition 3.1 identifies a single-threaded process of CLL with a single role implementation in MP, such that all its channels are mapped to the same multiparty session channel s and same principal role p , but to different destination roles.

CONVENTION 3.1. In the remainder of this section and Section 5, given $P \vdash_{\text{CL}}^{\sigma} \Delta$, we assume $\forall x, y \in \text{fn}(P), c_{\sigma}(x) = c_{\sigma}(y) = s$ and $p_{\sigma}(x) = p_{\sigma}(y) = p_{\sigma}$. This convention is allowed due to the session and principal role for all channels in a given mapping σ being constant. This assumption is lifted in Section 6.

Let $P \vdash_{\text{CL}}^{\sigma} \Delta$. We write $\sigma(P)$ for the process obtained by renaming each free name x in P with $\sigma(x)$, where actions in P are mapped to their corresponding actions in MP:

$$\begin{aligned}
\sigma(x(y).P) &\triangleq s[p_{\sigma}][d_{\sigma}(x)](y); \sigma(P) & \sigma(\bar{x}\langle y \rangle.P) &\triangleq s[p_{\sigma}][d_{\sigma}(x)]\langle y \rangle; \sigma(P) \\
\sigma(x.l_j; P) &\triangleq s[p_{\sigma}][d_{\sigma}(x)] \oplus l_j; \sigma(P) & \sigma(x.\text{case}\{l_i:Q_i\}_{i \in I}) &\triangleq s[p_{\sigma}][d_{\sigma}(x)] \& \{l_i:\sigma(Q_i)\}_{i \in I}.
\end{aligned}$$

Mapping 2: Generating Local Types. Having constructed a syntactic mapping from CLL to MP processes, we now present a way to generate the appropriate local typings for processes in the image of the translation.

Definition 3.2 (Local Type Generation). Let $P \vdash_{\text{CL}}^{\sigma} \Delta$. We generate a local type T such that $\sigma(P) \vdash_{\text{MP}} s[p_{\sigma}]:T$ by induction on the structure of P , written $\llbracket P \rrbracket_{\sigma}$ (assume $d_{\sigma}(x) = q$ and $S = \text{end}$, noting that value passing is encoded by the communication of sessions of type end):

$$\begin{aligned}
\llbracket 0 \rrbracket_{\sigma} &\triangleq \text{end} & \llbracket \bar{x}\langle y \rangle.P \rrbracket_{\sigma} &\triangleq q \uparrow(S); \llbracket P \rrbracket_{\sigma} & \llbracket x(y).P \rrbracket_{\sigma} &\triangleq q \downarrow(S); \llbracket P \rrbracket_{\sigma} \\
\llbracket x.l_j; P \rrbracket_{\sigma} &\triangleq \oplus q \{l_j: \llbracket P \rrbracket_{\sigma}\} & \llbracket x.\text{case}\{l_i:P_i\}_{i \in I} \rrbracket_{\sigma} &\triangleq \& q \{l_i: \llbracket P_i \rrbracket_{\sigma}\}_{i \in I}.
\end{aligned}$$

Hence, given a cut-free $P \vdash_{\text{CL}} \Delta$, we have an automatic way of generating a renaming σ such that $P \vdash_{\text{CL}}^{\sigma} \Delta$ and $\sigma(P) \vdash_{\text{MP}} s[p_{\sigma}]:T$ with $T = \llbracket P \rrbracket_{\sigma}$.

Mapping 3: Parallel Composition. Figure 4 defines the judgement $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$ such that P is an n -ary composition of processes, Γ is an MP session typing context, Δ is a CLL linear context, σ (respectively, ρ) is a mapping from free (respectively, bound) names to indexed channels. Recall that s stands for the (unique) channel in the mappings σ and σ' . Rule (comp) defines the composition of a single-thread CLL process with an n -ary composition of CLL processes that can be mapped to MP typed processes. The rule ensures that the resulting process is well formed in both CLL and MP: clause **(a – bound channel)** constructs the mapping ρ' for bound channels, as they are hidden by CLL composition; **(b – role/destination match)** ensures that σ and σ' map x to the same multiparty session channel, where the destination role in $\sigma(x)$ matches the principal role in $\sigma'(x)$, and vice versa; **(c – unique destinations)** asserts that channels in Δ and Δ' cannot have the same destination role, ensuring uniqueness of common channels and that free name assignments do not capture those of bound names.

We write $\rho(P)$ for the renaming of bound names in P generated by

$$\rho((\nu x)(P \mid Q)) = \rho'(P\{\rho(x)/x\}) \mid \rho'(Q\{\overline{\rho(x)}/x\}),$$

where $\rho' = \rho \setminus \{x\}$ and $\overline{\rho(x)}$ denotes $s[q][p]$ if $\rho(x) = s[p][q]$ (with the congruence cases).

Examples of the Translation. We give three examples of the translation from CLL to MP.

Example 3.3 (Conditions of Comp). We explain the conditions of comp via a small example. Consider the following processes:

$$\begin{aligned} P &\triangleq x\langle 7 \rangle . z\langle \text{“hello”} \rangle . \mathbf{0} & Q_1 &\triangleq z(u) . y(w) . \mathbf{0} & Q_2 &\triangleq z(v) . \mathbf{0} \mid y(w) . \mathbf{0} \\ P &\vdash_{\text{CL}} z:\text{str} \otimes \mathbf{1}, x:\text{nat} \otimes \mathbf{1} & Q_i &\vdash_{\text{CL}} z:\text{str} \wp \perp, y:\text{nat} \wp \perp. \end{aligned}$$

We define σ , σ_1 , and σ_2 such that:

$$\begin{aligned} \sigma(P) &= s[p][r]\langle 7 \rangle ; s[p][q]\langle \text{“hello”} \rangle \\ \sigma_1(Q_1) &= s[q][p](u) ; s[r_1][r_2](w) \\ \sigma_2(Q_2) &= s[q][p](v) \mid s[r_1][r_2](w). \end{aligned}$$

Then, assuming $\rho = \{\}$, the mappings σ and σ_i above satisfy **(b)** $(z, s[p][q]) \in \sigma$ and $(z, s[q][p]) \in \sigma_i$ by $p = p_\sigma(z) = d_{\sigma_i}(z)$ and $q = d_\sigma(z) = p_{\sigma_i}(z)$; **(c)** $(x, s[p][r]) \in \sigma$ and $(y, s[r_1][r_2]) \in \sigma_i$ with $r_1 = q$ in σ_1 and $r_1 \neq q$ in σ_2 ; and $d_\sigma(x) = r \neq r_2 = d_{\sigma_i}(y)$; and $r, r_2 \notin \rho$.

Example 3.4 (Four Threads). We show how to translate and compose the CLL process P from the beginning of this section:

$$\begin{aligned} P &\triangleq x\langle 7 \rangle . y(z) . x\langle \text{“hello”} \rangle . \mathbf{0} \vdash_{\text{CL}} x:\text{nat} \otimes \text{str} \otimes \mathbf{1}, y:\text{nat} \wp \perp \\ Q_1 &\triangleq x(x_1) . w\langle 93 \rangle . x(x_2) . \mathbf{0} \vdash_{\text{CL}} x:\text{nat} \wp \text{str} \wp \perp, w:\text{nat} \otimes \mathbf{1} \\ Q_2 &\triangleq y\langle 2 \rangle . \mathbf{0} \vdash_{\text{CL}} y:\text{nat} \otimes \mathbf{1} & Q_3 &\triangleq w(x_3) . \mathbf{0} \vdash_{\text{CL}} w:\text{nat} \wp \perp \end{aligned}$$

We define σ , σ_1 , σ_2 , σ_3 such that:

$$\begin{aligned} \sigma(P) &= s[p][q]\langle 7 \rangle ; s[p][r](z) ; s[p][q]\langle \text{“hello”} \rangle ; \mathbf{0} & \llbracket P \rrbracket_\sigma &= q \uparrow(\text{nat}) ; r \downarrow(\text{nat}) ; q \uparrow(\text{str}) ; \text{end} \\ \sigma_1(Q_1) &= s[q][p](x_1) ; s[q][s]\langle 93 \rangle ; s[q][p](x_2) ; \mathbf{0} & \llbracket Q_1 \rrbracket_{\sigma_1} &= p \downarrow(\text{nat}) ; s \uparrow(\text{nat}) ; p \downarrow(\text{str}) ; \text{end} \\ \sigma_2(Q_2) &= s[r][p]\langle 2 \rangle ; \mathbf{0} & \llbracket Q_2 \rrbracket_{\sigma_2} &= p \uparrow(\text{nat}) ; \text{end} \\ \sigma_3(Q_3) &= s[s][q](x_3) ; \mathbf{0} & \llbracket Q_3 \rrbracket_{\sigma_3} &= q \downarrow(\text{nat}) ; \text{end}. \end{aligned}$$

Let $\Gamma = s[p] : \llbracket P \rrbracket_\sigma, s[q] : \llbracket Q_1 \rrbracket_{\sigma_1}, s[r] : \llbracket Q_2 \rrbracket_{\sigma_2}, s[s] : \llbracket Q_3 \rrbracket_{\sigma_3}$. Then we have: $(\nu x, y, w)(P \mid Q_1 \mid Q_2 \mid Q_3) \Vdash_\rho^\emptyset ; \Gamma$.

Example 3.5 (Choice and Branching). As we discuss in Section 4, this CLL typable branching behaviour is not typable in the previous work on multiparty logic [11, 14] using the same local types:

$$\begin{aligned} P &\triangleq x.\text{case}\{l_1:y.l_2; \mathbf{0}, l_3:y.l_4; \mathbf{0}\} & Q_1 &\triangleq x.l_1; \mathbf{0} \\ R &\triangleq y.\text{case}\{l_2:\mathbf{0}, l_4:\mathbf{0}\} & Q_2 &\triangleq x.l_3; \mathbf{0} \end{aligned}$$

with $P \vdash_{\text{CLL}} x: \& \{l_1:\perp, l_3:\perp\}, y: \oplus \{l_2:1, l_4:1\}, R \vdash_{\text{CLL}} y: \& \{l_2:\perp, l_4:\perp\}$, and $Q_i \vdash_{\text{CLL}} x: \oplus \{l_1:1, l_3:1\}$. We define mappings σ, σ_1 , and σ_2 such that

$$\begin{aligned} \sigma(P) &= s[p][q] \& \{l_1:s[p][r] \oplus l_2; \mathbf{0}, l_3:s[p][r] \oplus l_4; \mathbf{0}\} \\ \sigma_1(Q_1) &= s[q][p] \oplus l_1; \mathbf{0} \\ \sigma_1(Q_2) &= s[q][p] \oplus l_3; \mathbf{0} \\ \sigma_2(R) &= s[r][p] \& \{l_2;\mathbf{0}, l_4;\mathbf{0}\} \\ \llbracket P \rrbracket_\sigma &= \&q\{l_1: \oplus r\{l_2:\text{end}\}, l_3: \oplus r\{l_4:\text{end}\}\} \\ \llbracket Q_1 \rrbracket_{\sigma_1} &= \oplus p\{l_1:\text{end}\} \\ \llbracket Q_2 \rrbracket_{\sigma_1} &= \oplus p\{l_3:\text{end}\} \\ \llbracket R \rrbracket_{\sigma_2} &= \&p\{l_2:\text{end}, l_4:\text{end}\}. \end{aligned}$$

Let $\Gamma = s[p]:\llbracket P \rrbracket_\sigma, s[q]:\llbracket Q_i \rrbracket_{\sigma_1}, s[r]:\llbracket R \rrbracket_{\sigma_2}$. Thus, we have $(\nu x, y)(P \mid Q_i \mid R) \Vdash_{\rho}^{\emptyset} \cdot; \Gamma$.

Type-preservation and Uniqueness. Below we study properties of the encoding. We first show that the type-preserving translation of CLL to MP for cut-free processes combined with our composition rule preserves typing in MP.

PROPOSITION 3.6 (TYPE PRESERVATION). *If $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$, then $\rho(\sigma(P)) \vdash_{\text{MP}} \Gamma$.*

PROOF. The prefix case is straightforward by Definition 3.1 and (thread); the parallel composition uses (comp). Both cases are mechanical by induction on P . $\square$

Since the mapping from CLL into MP is just renaming, reduction of CLL strongly corresponds to that of MP.

PROPOSITION 3.7 (OPERATIONAL CORRESPONDENCE). *Suppose $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$ and $P \rightarrow P'$. Then $\rho(\sigma(P)) \rightarrow Q$ s.t. $P' \Vdash_{\rho'}^{\sigma'} \Delta'; \Gamma'$ and $Q = \rho'(\sigma'(P'))$ with $\sigma' \subseteq \sigma, \rho' \subseteq \rho$.*

PROOF. See Appendix A.1.1. $\square$

We call a mapping *thread preserving* if it assigns to a cut-free CLL process a single participant in MP. We thus have the following:

PROPOSITION 3.8 (THREAD PRESERVATION). *If $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$, then $\rho(\sigma(P))$ is thread preserving.*

PROOF. See Appendix A.1.1. $\square$

Theorem 3.9 states that the mapping is closed under any bijective renaming φ on sessions, roles, and channels. As an example, let $P = s[p][r](x); s[p][q]\langle v \rangle$ and $\varphi = \{s \mapsto s', x \mapsto y, p \mapsto p'\}$; then $\varphi(P) = s'[p']\langle r \rangle(y); s'[p']\langle q \rangle\langle v \rangle$.

More precisely, Theorem 3.9 shows that any thread-preserving mapping from CLL processes into a single MP session always conforms to our mapping. This means that no other way to encode CLL into MP (modulo bijective renaming that maps different names to distinct destinations) exists if it is thread preserving into a single multiparty session.

THEOREM 3.9 (UNIQUENESS). *Assume $P \vdash_{\text{CLL}} \Delta$. Suppose $\varphi(P)$ is thread preserving and $\varphi(P)$ is typable by a single MP session, i.e., if $\varphi(P) \vdash_{\text{MP}} \Gamma$ then (1) $\text{dom}(\Gamma)$ contains a single session channel; or (2) $\Gamma = \emptyset$ and $P \equiv \mathbf{0}$. Then there exist ρ and σ such that $\varphi = \sigma \circ \rho$ and $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$.*

PROOF. See Appendix A.1. $\square$

4 PARTIAL MULTIPARTY COMPATIBILITY

This section studies a compositional synthesis property, dubbed PMC. As illustrated in Proposition 2.7, multiparty session type theories [31] cannot ensure deadlock-freedom if we do not rely on (1) a projection from a global type or (2) a global synthesis property called multiparty compatibility [22, 37], which given local types for a *complete* MP session produces a global type if the endpoints do not deadlock. For example, the previous counterexample can be avoided if we start from the global type (again, we assume basic value passing, encodable with terminated sessions):

$$G = p \rightarrow q:(\text{nat}).q \rightarrow r:(\text{bool}).r \rightarrow p:(\text{str}).\text{end} \quad (2)$$

and type each process with the projected local types:

$$T_p = q \uparrow (\text{nat}); r \downarrow (\text{str}); \text{end}, \quad T_q = p \downarrow (\text{nat}); r \uparrow (\text{bool}); \text{end}, \quad T_r = q \downarrow (\text{bool}); p \uparrow (\text{str}); \text{end} \quad (3)$$

or we may build (synthesise) G in (2) from $\{T_p, T_q, T_r\}$ in Equation (3). If we start from a projectable global type or can synthesise a global type, then the example in Proposition 2.7 is no longer typable.

Given that CLL employs a binary form of composition, we move from a global synthesis condition to a binary (partial) relation to achieve our main results. Specifically, we take the following steps:

- Step 1:** We introduce *partial* global types $p \rightsquigarrow q$ representing global interaction that has not yet been composed with another party (e.g., it denotes the emission from p to q , not yet composed with the reception by q), and give formal semantics to both global and local types (Section 4.1) as labelled transition systems. Crucially, the semantics of global types is given up to a swapping relation $\sim_{\text{sw}}$, which enables the permutation of independent actions.
- Step 2:** We define synchronous multiparty compatibility (SMC, Definition 4.7), showing the equivalence of SMC, deadlock-freedom and the existence of a global type that corresponds to the appropriate local behaviours.
- Step 3:** We introduce a notion of fusion (Definition 4.10), which enables us to compose compatible partial specifications and define *partial* multiparty compatibility. When a $p \rightsquigarrow q$ arrow denoting a send action is fused with the corresponding arrow denoting the receive action, it is transformed into a complete arrow $p \rightarrow q$, preserving the ordering of communications. When we compose all participants in a session (reconstructing a *complete* global type—one without partial arrows), deadlock-freedom is guaranteed (Theorem 4.15).

4.1 Partial Global Types and Semantics

We define partial global types G , consisting of a combination of complete global types and endpoint interactions.

Definition 4.1 (Partial Global Types). The grammar of partial global types G, G' is

$$\begin{aligned} G ::= & \text{end} \mid p \rightarrow q:(T).G \mid p \rightarrow q:\{l_j:G_j\}_{j \in J} \\ & \mid p \rightsquigarrow q:\uparrow(T).G \mid p \rightsquigarrow q:\downarrow(T).G \mid p \rightsquigarrow q:\oplus\{l_j:G_j\}_{j \in J} \mid p \rightsquigarrow q:\&\{l_j:G_j\}_{j \in J}. \end{aligned}$$

The first three of the above grammar constructs are the standard global types [31]. Global type $p \rightarrow q:(T).G$ means that participant p sends a session endpoint of type T to participant q , followed by G . Global type $p \rightarrow q:\{l_j:G_j\}_{j \in J}$ means that participant p selects label l_i , then q 's i th branch will be chosen, becoming G_i . The *partial* global types in the second line denote *half* of a complete global interaction. The modes ($\uparrow, \downarrow, \oplus, \&$) in partial global types indicate which component of the interaction is being satisfied: e.g., $p \rightsquigarrow q:\uparrow(T)$ denotes the contribution of the emission component of the interaction from principal p to destination q , whereas $p \rightsquigarrow q:\downarrow(T)$ denotes the reception.

We write mode $\dagger$ for either $\uparrow, \downarrow, \oplus, \&$ or $\emptyset$ (empty) and often omit $\dagger$ from partial global types when unnecessary. We write $\rightarrow$ for either $\rightarrow$ or $\leadsto$; and $p \leftrightarrow q$ for either $p \rightarrow q$ or $q \rightarrow p$.

The set of *principal roles* is defined as: $\text{pr}(p \rightarrow q : \emptyset) = \{p, q\}$ and $\text{pr}(p \leadsto q : \uparrow) = \text{pr}(q \leadsto p : \downarrow) = \text{pr}(p \leadsto q : \oplus) = \text{pr}(q \leadsto p : \&) = \{p\}$. We write $\text{roles}(G)/\text{pr}(G)$ for the set of roles/principal roles occurring in G ; and $p \leftrightarrow q \in G$ if $p \leftrightarrow q$ occurs in G .

We use standard projection rules from global to local types, defined in terms of a merge operation for branchings [22], written $T \sqcup T'$, ensuring that if the locally observable behaviour of the local type is not independent of the chosen branch then it is identifiable via a unique label (the operator is otherwise undefined).

Definition 4.2 (Local Type Merge). The merge $T \sqcup T'$ of T and T' is defined by $T \sqcup T \triangleq T$; and with $T = \&r\{l_i : T_i\}_{i \in I}$ and $T' = \&r\{l'_j : T'_j\}_{j \in J}$,

$$T \sqcup T' \triangleq \&r(\{l_h : T_h\}_{h \in I \cap J} \cup \{l'_h : T'_h\}_{h \in J \setminus I} \cup \{l_h : T_h \sqcup T'_h\}_{h \in I \cap J})$$

if $l_h = l'_h$ for each $h \in I \cap J$, and homomorphic for other types (i.e., $\mathcal{T}[T_1] \sqcup \mathcal{T}[T_2] = \mathcal{T}[T_1 \sqcup T_2]$, where $\mathcal{T}$ is a context of local types). $T \sqcup T'$ is undefined otherwise.

Definition 4.3 (Projection and Well-formedness). Let G be a global type. The projection of G for a role p , written $G \upharpoonright p$, is defined below.

$$\begin{aligned} \text{end} \upharpoonright p &= \text{end} \\ s \rightarrow r(T).G' \upharpoonright p &= \begin{cases} r \uparrow(T); (G' \upharpoonright p) & \text{if } p = s \\ s \downarrow(T); (G' \upharpoonright p) & \text{if } p = r \\ G' \upharpoonright p & \text{otherwise} \end{cases} \\ s \rightarrow r\{l_j : G_j\}_{j \in J} \upharpoonright p &= \begin{cases} \oplus r\{l_j : G_j \upharpoonright p\}_{j \in J} & \text{if } p = s \\ \&s\{l_j : G_j \upharpoonright p\}_{j \in J} & \text{if } p = r \\ \sqcup_{j \in J} G_j \upharpoonright p & \text{otherwise.} \end{cases} \end{aligned}$$

If no side conditions hold (i.e., the merge operator is undefined), then projection is undefined. We say that G is *well formed* iff for all distinct $p \in \text{roles}(G)$, $(G \upharpoonright p)$ is defined.

As an illustration of merge and projection, consider

$$G = q \rightarrow p : \{l_1 : p \rightarrow r : \{l_2 : \text{end}\}, l_3 : p \rightarrow r : \{l_4 : \text{end}\}\},$$

which will be built from CLL processes in Example 3.5. Then:

$$\begin{aligned} G \upharpoonright p &= \&q\{l_1 : \oplus r\{l_2 : \text{end}\}, l_3 : \oplus r\{l_4 : \text{end}\}\}, \\ G \upharpoonright q &= \oplus p\{l_1 : \text{end}, l_3 : \text{end}\} \text{ and } G \upharpoonright r = \&p\{l_2 : \text{end}, l_4 : \text{end}\}. \end{aligned} \quad (4)$$

The syntax of global types can impose unnecessary orderings of actions among independent roles. For example, $p \rightarrow q : (\text{str}).r \rightarrow s : (\text{bool}).\text{end}$ should be regarded as identical to $r \rightarrow s : (\text{bool}).p \rightarrow q : (\text{str}).\text{end}$ if p and q do not coincide with either r or s , since there is no reasonably enforceable ordering between the two interactions. Thus, we allow for the swapping of independent communication actions in global types as defined below (a similar swapping relation is used [14] to define coherence of logical global types and in the context of semantics for choreographies [12]).

Definition 4.4 (Swapping). We define the swapping relation, written $\sim_{\text{sw}}$, as the smallest congruence on global types satisfying (with $\text{pr}(p \rightarrow q : \dagger) \cap \text{pr}(p' \rightarrow q' : \dagger) = \emptyset$):

- (ss) $p \rightarrow q : \dagger(T).p' \rightarrow q' : \dagger'(T').G \sim_{\text{sw}} p' \rightarrow q' : \dagger'(T').p \rightarrow q : \dagger(T).G$
- (sb) $p \rightarrow q : \dagger(T).p' \rightarrow q' : \dagger'\{l_i : G_i\}_{i \in I} \sim_{\text{sw}} p' \rightarrow q' : \dagger'\{l_i : p \rightarrow q : \dagger(T).G_i\}_{i \in I}$
- (bb) $p \rightarrow q : \dagger\{l_i : p' \rightarrow q' : \dagger'\{l'_j : G_j\}_{j \in J}\}_{i \in I} \sim_{\text{sw}} p' \rightarrow q' : \dagger'\{l'_j : p \rightarrow q : \dagger\{l_i : G_j\}_{i \in I}\}_{j \in J}$.

Local Type (at role p)

$$q \uparrow (S); T \xrightarrow{pq \uparrow (S)} T \quad q \downarrow (S); T \xrightarrow{pq \downarrow (S)} T \quad \oplus q\{l_i : T_i\}_{i \in I} \xrightarrow{pq \triangleleft l_k} T_k \quad (k \in I) \quad \& q\{l_i : T_i\}_{i \in I} \xrightarrow{pq \triangleright l_k} T_k \quad (k \in I)$$

Global Type

$$p \rightarrow q; (T).G \xrightarrow{pq \uparrow (T) \cdot qp \downarrow (T)} G \quad p \rightarrow q; \{l_i : G_i\}_{i \in I} \xrightarrow{pq \triangleleft l_k \cdot qp \triangleright l_k} G_k \quad (k \in I)$$

$$G_1 \sim_{sw} G'_1 \wedge G'_1 \xrightarrow{\ell \cdot \ell'} G'_2 \wedge G'_2 \sim_{sw} G_2 \Rightarrow G_1 \xrightarrow{\ell \cdot \ell'} G_2$$

Configuration

$$(T_p \xrightarrow{\ell} T'_p) \wedge (T_q \xrightarrow{\bar{\ell}} T'_q) \wedge (\forall r \in \mathcal{P} \setminus \{p, q\}. T_r = T'_r) \Rightarrow (T_p)_{p \in \mathcal{P}} \xrightarrow{\ell \cdot \bar{\ell}} (T'_p)_{p \in \mathcal{P}} \quad (\ell \text{ output or selection})$$

Fig. 5. Labelled transition systems.

The operational semantics for local, global types and configurations are given by labelled transition systems (LTS). We define the syntax of labels as:

$$\ell ::= pq \uparrow (T) \mid pq \downarrow (T) \mid pq \triangleleft l \mid pq \triangleright l,$$

where $pq \uparrow (T)$ (respectively, $pq \downarrow (T)$) is the output (respectively, input) at p to q (respectively, from q) and $pq \triangleleft l$ (respectively, $pq \triangleright l$) is the selection (respectively, branching).

We define $\bar{\ell}$ as $\bar{pq} \uparrow (\tau) = qp \downarrow (\tau)$ and $\bar{pq} \triangleleft l = qp \triangleright l$ and vice versa. Given a set of roles $\mathcal{P}$, we define a *configuration* as $C = (T_p)_{p \in \mathcal{P}}$. Configurations consist of a set of local types projected from a single global type that are used to define and show properties of local types in Section 4.2.

Definition 4.5 (Labelled Transition Relations). Transitions between local types, written $T \xrightarrow{\ell} T'$, for role p ; global types, written $G \xrightarrow{\ell \cdot \ell'} G'$; and configurations are given in Figure 5. We write $G \xrightarrow{\bar{\ell}}$ G_n if $G \xrightarrow{\ell_1 \cdot \ell_2} G_1 \cdots \xrightarrow{\ell_{2n-1} \cdot \ell_{2n}} G_n$ and $\bar{\ell} = \ell_1 \cdots \ell_{2n}$ ($n \geq 0$); and $Tr(G_0) = \{\bar{\ell} \mid G_0 \xrightarrow{\bar{\ell}} G_n \ n \geq 0\}$ for *traces* of type G_0 . Similarly for T and C .

PROPOSITION 4.6 (TRACE EQUIVALENCE). Suppose G is well formed and the set of participants in G is $\mathcal{P}$. Assume $C = (G \upharpoonright p)_{p \in \mathcal{P}}$. Then $Tr(C) = Tr(G)$.

PROOF. By definition of the projection and the LTSs. $\square$

4.2 Partial Multiparty Compatibility

To introduce PMC, we first define multiparty compatibility (MC) for a *synchronous* semantics, adapting the development of MC for asynchrony [22, 37]. We then introduce PMC as a compositional binary synthesis property on types. We note that while asynchronous formulations of the linear logic-based session calculi exist [23], the predominant formulation is synchronous, and so we focus on a synchronous theory.

Definition 4.7 (Synchronous Multiparty Compatibility). Configuration $C_0 = (T_{0p})_{p \in \mathcal{P}}$ is *synchronous multiparty compatible* (SMC) if for all $C_0 \xrightarrow{\bar{\ell}} C = (T_p)_{p \in \mathcal{P}}$ and $T_p \xrightarrow{\ell} T'_p$:

- (1) if $\ell = pq \uparrow (S)$ or $pq \triangleleft l$, there exists $C \xrightarrow{\bar{\ell}'} C' \xrightarrow{\ell \cdot \bar{\ell}'} C''$;
- (2) if $\ell = pq \downarrow (S)$, there exists $C \xrightarrow{\bar{\ell}'} C' \xrightarrow{\bar{\ell} \cdot \ell} C''$; or
- (3) if $\ell = pq \triangleright l$, there exists $\ell_1 = pq \triangleright l'$, $C \xrightarrow{\bar{\ell}'} C' \xrightarrow{\ell_1 \cdot \bar{\ell}_1} C''$,

where $\bar{\ell}'$ does not include actions from or to p .

Since our semantics is synchronous, it is technically simpler than one for an asynchronous semantics [22, 37]. One can check that the local types in Equation (3) and the projected local types of (4.1) satisfy SMC.

Definition 4.8 (Deadlock-freedom). $C = (T_p)_{p \in \mathcal{P}}$ is *deadlock-free* if for all $C \xrightarrow{\vec{e}} C_1$, $\exists C' = (T'_p)_{p \in \mathcal{P}}$ such that $C_1 \xrightarrow{e'} C'$ or $T'_p = \text{end}$ for all $p \in \mathcal{P}$.

THEOREM 4.9 (DEADLOCK-FREEDOM, MC AND EXISTENCE OF A GLOBAL TYPE). *The following are equivalent: (MC) a configuration C is SMC; (DF) C is deadlock-free; (WF) there exists well-formed G such that $\text{Tr}(G) = \text{Tr}(C)$.*

PROOF. See Appendix A.2.1. □

Multiparty compatibility is a global property defined using the set of all participants [22, 37]. To define a compositional (i.e., local) multiparty compatibility, we introduce the composition of two partial global types, dubbed as *fusion*.

Definition 4.10 (Fusion). We define the fusion of two well-formed partial global types G_1, G_2 such that $\text{pr}(G_1) \cap \text{pr}(G_2) = \emptyset$, written $\text{fuse}(G_1, G_2)$, inductively on the structure of G_1 and G_2 , up to the swapping relation $\sim_{\text{sw}}$:

$$\begin{aligned} \text{fuse}(p \rightsquigarrow q : \uparrow(T_1).G'_1, p \rightsquigarrow q : \downarrow(T_2).G'_2) &= p \rightarrow q : (T_2).\text{fuse}(G'_1, G'_2) \quad (\text{with } T_1 \geq T_2) \\ \text{fuse}(p \rightsquigarrow q : \oplus\{l : G'_1\}, p \rightsquigarrow q : \&\{l : G'_2, \{l_j : G_j\}_{j \in J}\}) &= p \rightarrow q : \{l : \text{fuse}(G'_1, G'_2)\} \\ \text{fuse}(p \rightsquigarrow q : \dagger(T).G_1, G_2) &= p \rightsquigarrow q : \dagger(T).\text{fuse}(G_1, G_2) \\ \text{fuse}(p \rightsquigarrow q : \dagger\{l_j : G_j\}_{j \in J}, G_2) &= p \rightsquigarrow q : \dagger\{l_j : \text{fuse}(G_j, G_2)\}_{j \in J} \\ \text{fuse}(p \rightsquigarrow q : \dagger\{l_j : G_j\}_{j \in J}, G_2) &= p \rightsquigarrow q : \dagger\{l_j : G'_j\}_{j \in J} \sim_{\text{sw}} G_2 \wedge p \leftrightarrow q \notin G_2 \\ \text{fuse}(p \rightarrow q : (T).G_1, G_2) &= p \rightarrow q : (T).\text{fuse}(G_1, G_2) \text{ if } p \leftrightarrow q \notin G_2 \\ \text{fuse}(p \rightarrow q : \{l_j : G_j\}_{j \in J}, G_2) &= p \rightarrow q : \{l_j : \text{fuse}(G_j, G_2)\}_{j \in J} \text{ if } p \leftrightarrow q \notin G_2 \\ \text{fuse}(\text{end}, \text{end}) &= \text{end} \end{aligned}$$

with the symmetric cases.

The first rule uses the subtyping relation $T \leq S$ given in Definition 2.5. The second rule selects one branch with the same label. For simplicity, we allow only for one-way selections since in the context of our work, partial types are extracted from processes where the selections are always determined. We note that multi-way branchings can be realised straightforwardly via subtyping.

The third and fourth rules (which do not overlap with the first two) push through actions that are unmatched in the fused types. The rule is extended similarly to input, branching and selection with other global type constructors. $\text{fuse}(G_1, G_2)$ is undefined if none of the above rules are applicable.

We define (1) $|\text{end}| = 1$; (2) $|p \rightarrow q : \dagger(T).G| = 4 + |G|$; and (3) $|p \rightarrow q : \dagger\{l_j : G_j\}_{j \in J}| = 3 + \sum_{j \in J} (1 + |G_j|)$. We have the following:

PROPOSITION 4.11. *Computing $\text{fuse}(G_1, G_2)$ is $\mathcal{O}(|G_1|! \times |G_2|!)$ time in the worst case, where $|G|$ is the size of G .*

PROOF. The time complexity is dominated by the computation of the swapping relation between G_1 and G_2 . The equivalence class up to the swapping relation consists of $|G|!$ elements. Since we apply each fuse rule to the equivalence class of G_1 and the equivalence class of G_2 (in the worst case), the time complexity is $\mathcal{O}(|G_1|! \times |G_2|!)$. □

Definition 4.12 (Partial Multiparty Compatibility). Suppose G_1 and G_2 are partial global types. G_1 and G_2 are partial multiparty compatible iff $\text{fuse}(G_1, G_2)$ is defined.

Example 4.13. Consider the following partial global types for Example 3.5 (as mentioned before, the corresponding local types are not coherent in Carbone et al. [11, 14]):

$$\begin{aligned} G_1 &= q \rightsquigarrow p: \& \{l_1: p \rightsquigarrow r: \oplus \{l_2: \text{end}\}, l_3: p \rightsquigarrow r: \oplus \{l_4: \text{end}\}\} \\ G_2 &= q \rightsquigarrow p: \oplus \{l_1: \text{end}\} \\ G_3 &= p \rightsquigarrow r: \& \{l_2: \text{end}, l_4: \text{end}\} \\ G_4 &= q \rightsquigarrow p: \oplus \{l_3: \text{end}\} \end{aligned}$$

Then we have the following:

$$\begin{aligned} \text{fuse}(G_1, G_2) &= q \rightarrow p: \{l_1: p \rightsquigarrow r: \oplus \{l_2: \text{end}\}\} \\ \text{fuse}(G_2, G_3) &= q \rightsquigarrow p: \oplus \{l_1: G_3\} \\ \text{fuse}(\text{fuse}(G_1, G_2), G_3) &= \text{fuse}(G_1, \text{fuse}(G_2, G_3)) = q \rightarrow p: \{l_1: p \rightarrow r: \{l_2: \text{end}\}\} \\ \text{fuse}(G_1, G_4) &= q \rightarrow p: \{l_3: p \rightsquigarrow r: \oplus \{l_4: \text{end}\}\} \\ \text{fuse}(G_4, G_3) &= q \rightsquigarrow p: \oplus \{l_3: G_3\} \\ \text{fuse}(\text{fuse}(G_1, G_4), G_3) &= \text{fuse}(G_1, \text{fuse}(G_4, G_3)) = q \rightarrow p: \{l_3: p \rightarrow r: \{l_4: \text{end}\}\}. \end{aligned}$$

LEMMA 4.14. Suppose $\text{fuse}(\text{fuse}(G_i, G_j), G_k)$ with $\{i, j, k\} = \{1, 2, 3\}$ is well formed. Then we have $\text{fuse}(\text{fuse}(G_i, G_j), G_k) \sim_{\text{sw}} \text{fuse}(G_i, \text{fuse}(G_j, G_k))$.

PROOF. See Appendix A.2.2. □

By the above lemma, we have the following:

THEOREM 4.15 (COMPOSITIONALITY). Suppose $G_1, \dots, G_n$ are partial global types. Assume $\forall i, j$ such that $1 \leq i \neq j \leq n$, G_i and G_j are PMC and $G = \text{fuse}(G_1, \text{fuse}(G_2, \text{fuse}(\dots, G_n)))$ is a complete global type. Then G is well formed.

PROOF. See Appendix A.2.3. □

5 ENCODING CLL AS A SINGLE MULTIPARTY SESSION

Having defined in Section 3 how to translate CLL processes to MP, we study the interconnection networks induced by CLL by generating their partial global types (Section 4). We note that such types are well formed by construction. We prove a strict inclusion of the networks of CLL into those of single MP, by fusing the partial global types into a *complete* global type.

Definition 5.1 (Generating Partial Global Types). Given P with $P \vdash_{\text{CL}}^{\sigma} \Delta$, we generate its partial global type, written $\langle P \rangle_{\sigma}$ as follows (let $d_{\sigma}(x) = q$):

$$\begin{aligned} \langle 0 \rangle_{\sigma} &\triangleq \text{end} \quad \langle \bar{x}(y).P \rangle_{\sigma} \triangleq p_{\sigma} \rightsquigarrow q: \uparrow(T). \langle P \rangle_{\sigma} \quad \langle x(y).P \rangle_{\sigma} \triangleq q \rightsquigarrow p_{\sigma}: \downarrow(T). \langle P \rangle_{\sigma} \\ \langle x.l; P \rangle_{\sigma} &\triangleq p_{\sigma} \rightsquigarrow q: \oplus \{l: \langle P \rangle_{\sigma}\} \quad \langle x.\text{case}\{l_i : P_i\}_{i \in I} \rangle_{\sigma} \triangleq q \rightsquigarrow p_{\sigma}: \& \{l_i: \langle P_i \rangle_{\sigma}\}_{i \in I}. \end{aligned}$$

We generate a set $\mathcal{G}$ of partial global types for compositions, written $P \Vdash_{\rho}^{\sigma} \Delta; \mathcal{G}$, by (we omit the obvious (thread) rule):

$$\frac{(\text{comp-}\mathcal{G}) \quad P \vdash_{\text{CL}}^{\sigma} \Delta, x:A \quad Q \Vdash_{\rho}^{\sigma'} \Delta', x:A^{\perp}; \mathcal{G} \quad (\star) \text{ in (comp)}}{(vx)(P \mid Q) \Vdash_{\rho'}^{(\sigma' \cup \sigma) \setminus \{a\}} \Delta, \Delta'; \mathcal{G} \cup \langle P \rangle_{\sigma}}.$$

Example 5.2 (Four Threads). We present the generated partial global types for Example 3.4:

$$\begin{aligned} \langle P \rangle_{\sigma} &= p \rightsquigarrow q: \uparrow(\text{nat}).r \rightsquigarrow p: \downarrow(\text{nat}).p \rightsquigarrow q: \uparrow(\text{str}).\text{end} \\ \langle Q \rangle_{\sigma_2} &= p \rightsquigarrow q: \downarrow(\text{nat}).q \rightsquigarrow s: \uparrow(\text{nat}).p \rightsquigarrow q: \downarrow(\text{str}).\text{end} \\ \langle R \rangle_{\sigma_1} &= r \rightsquigarrow p: \uparrow(\text{nat}).\text{end} \quad \langle S \rangle_{\sigma_3} = q \rightsquigarrow s: \downarrow(\text{nat}).\text{end}. \end{aligned}$$

Where applying fuse to the partial global types above produces the global type: $p \rightarrow q:(\text{nat}).q \rightarrow s:(\text{nat}).r \rightarrow p:(\text{nat}).p \rightarrow q:(\text{str}).\text{end}$. We note that, for instance, adding a message from r to s makes the example untypable in CLL since it introduces a three-way cycle in the interconnection network.

Example 5.3 (Choice and Branching). The partial global types generated for Example 3.5 are $G_1 = \langle P \rangle_\sigma$, $G_2 = \langle Q_1 \rangle_{\sigma_1}$, $G_3 = \langle R \rangle_{\sigma_2}$, and $G_4 = \langle Q_2 \rangle_{\sigma_1}$ in Example 4.13. They fuse into: $q \rightarrow p:\{l_1:p \rightarrow r:\{l_2:\text{end}\}\}$ or $q \rightarrow p:\{l_3:p \rightarrow r:\{l_4:\text{end}\}\}$. Notably, the global type produced by our example cannot be captured by Carbone et al. [11, 14], requiring the two selections in branches l_1 and l_3 to be the same.

We make precise the claims of Section 1, that is, the network interconnections of global types induced by CLL are strictly less expressive than those of MP, by formalising the notion of interconnection network as an undirected graph.

Definition 5.4 (Interconnection Network). Given a partial global type G we generate its *Interconnection Network Graph* (ING), where the nodes are the roles of G and two nodes p, q share an edge iff $p \leftrightarrow q$ in G .

We establish the main properties of our framework. The following proposition states that we can always fuse the partial global types of a well-formed composition.

PROPOSITION 5.5. *Let $P \Vdash_\rho^\sigma \Delta; \mathcal{G}$ and $\Delta = \emptyset$ or Δ contains only $\mathbf{1}$ or $\perp$. There exists a single well-formed global type G such that $G = \text{fuse}(\mathcal{G})$, where $\text{fuse}(\mathcal{G})$ denotes fusion of all partial global types in $\mathcal{G}$.*

PROOF. Since Δ is empty or contains only $\mathbf{1}$ or $\perp$, we have that P is an n -ary composition of (cut-free) processes. If $n = 1$, then $P = \mathbf{0}$ and its corresponding global type is just end . The interesting case is when $n > 1$.

Since the context is either empty or contains only $\mathbf{1}$ or $\perp$, we have that P is of the form $(\nu \tilde{a})(P_1 \mid \dots \mid P_n)$ where all free names $a_j:A_j$ of each of the P_i processes are cut with some other $P_{i'}$ using $a_j:A_j^\perp$. Thus, by construction of $\Vdash$, we have that for each bound name a of P we have $c_\rho(a)[p_\rho(a)]:T \in \Gamma$ and $c_\rho(a)[d_\rho(a)]:T' \in \Gamma$ with $T \upharpoonright d_\rho(a) \leq \overline{T' \upharpoonright p_\rho(a)}$ and thus for each action between two roles in a partial global type in $\mathcal{G}$, we can always find a matching action in another partial global type in $\mathcal{G}$; therefore, we can fuse all partial global types in $\mathcal{G}$ into a single global type. $\square$

The following main theorem shows that any two such partial types overlap in at most 2 roles; from which the acyclicity of CLL processes follows, establishing a separation between MP global types and those induced by CLL. Notice that Theorem 3.9 ensures that, in general, there exists no type- and thread-preserving translation from CLL to a single MP session where CLL has the same or more interconnectability than MP.

THEOREM 5.6. *Let $P \Vdash_\rho^\sigma \Delta; \mathcal{G}$. For any distinct $G_1, G_2 \in \mathcal{G}$, we have that $\text{roles}(G_1) \cap \text{roles}(G_2)$ contains at most 2 elements.*

PROOF. We proceed by induction on the derivation $P \Vdash_\rho^\sigma \Delta; \mathcal{G}$, showing that each case preserves the specified invariant of at most two elements in the intersection of $\text{roles}(G_1) \cap \text{roles}(G_2)$, for any $G_1, G_2 \in \mathcal{G}$.

The only interesting case is when the last rule in the derivation is $(\text{comp-}\mathcal{G})$:

$$\frac{(\text{comp-}\mathcal{G}) \quad \begin{array}{c} P \Vdash_{\text{CL}}^\sigma \Delta, x:A \quad Q \Vdash_{\rho'}^{\sigma'} \Delta', x:A^\perp; \mathcal{G} \quad (\star) \text{ in } (\text{comp}) \end{array}}{(\nu x)(P \mid Q) \Vdash_{\rho'}^{(\sigma' \cup \sigma) \setminus \{x\}} \Delta, \Delta'; \mathcal{G} \cup \{\langle P \rangle_\sigma\}}.$$

By the inductive hypothesis, we have that for any $G'_1, G'_2 \in \mathcal{G}$, $\text{roles}(G'_1) \cap \text{roles}(G'_2)$ contains at most two elements.

By construction, we know that roles in $\mathcal{G}$ must either appear in σ' (corresponding to role assignments to channels in Δ' and a) or ρ (corresponding to role assignments to bound names).

By inversion, we know that $\forall z \in \Delta, y \in \Delta'. d_\sigma(z) \neq d_{\sigma'}(y)$, thus there are no common d_σ role assignments between σ and σ' to free names of the two processes beyond those for x . We also know that $p_\sigma(x) = d_{\sigma'}(x)$ and $d_\sigma(x) = p_{\sigma'}(x)$.

By the definition of $\langle P \rangle_\sigma$ there are at least two common role names with each endpoint interaction in $\mathcal{G}$ coming from σ and σ' (i.e., role assignments to free names), which are p_σ and $p_{\sigma'}$. Since p_σ is invariant and $\forall z \in \Delta, y \in \Delta'. d_\sigma(z) \neq d_{\sigma'}(y)$, we have that free names in Δ cannot share any additional roles.

We need now only consider ρ . By construction, we know that $\forall z \in \Delta, y \in \Delta'. d_\sigma(z) \notin \rho \wedge d_{\sigma'}(y) \notin \rho$, thus $d_\sigma(z)$ cannot appear in $\mathcal{G}$ due to ρ . The only remaining possibilities are $p_\sigma(x)$ and $d_\sigma(x)$ that are already accounted for from the argument above. Thus we preserve the invariant and conclude the proof. $\square$

We can then establish a main result of our work: The connection network graphs generated by CLL-typable processes are acyclic.

THEOREM 5.7. *Let $P \Vdash_\rho^\sigma \Delta; \mathcal{G}$. Let $G = \text{fuse}(\mathcal{G})$. The interconnection network graph for G is acyclic.*

PROOF. Each endpoint interaction sequence in $\mathcal{G}$ denotes the contribution of a single endpoint role in the global conversation. By Theorem 5.6, we have that any distinct pair of partial global types in $\mathcal{G}$ shares at most 2 role names. This means that for any distinct roles p, q, r , if $p \leftrightarrow q \in G$ and $p \leftrightarrow r \in G$ then neither $q \rightarrow r$ nor $r \rightarrow q$ or $q \rightsquigarrow r$ nor $r \rightsquigarrow q$ in G . Hence, in the connection graph of G we know that we cannot have triangles of the form $(p, q), (p, r), (r, q)$ as edges.

We can then see that no cycles can be formed through a “diamond”—a sequence of edges of the form $(p, q), (p, r), (q, t_1), \dots, (t_n, s), (r, v_1), \dots, (v_m, s)$ —in the graph by the fact that at each composition step, processes can only share one free name (the one that is the focus of the composition rule since $\Delta \cup \Delta' = \emptyset$) and role assignments ($\forall z \in \Delta, y \in \Delta'. d_\sigma(z) \neq d_{\sigma'}(y)$, similarly for ρ , by $(\star)$ (c) in rule (comp) and Definition 3.1). If we could form a “diamond” cycle in the graph, then we have to be able to eventually compose processes sharing more than one name or with different roles mapping to the same channel name to connect both (v_m, s) and (t_n, s) . That is, after composing v_m we cannot compose the implementation of t_n (or vice versa), since it would violate the role assignment restriction of composition— $(\star)$ (c) in (comp)—which disallows identical destination roles. Moreover, to compose t_n and s it would have to be the case that s shares a channel with both t_n and v_m (which is untypable) or the process composition of implementations of v_m and t_n would have to map both roles to the same channel shared with the implementation of s —itself also a contradiction. $\square$

Deadlock-freedom in MP. Theorem 5.10 states that our encoding produces a single multiparty session, that is, the fusing of all partial global types in a complete session is deadlock-free. To prove the theorem, we require the following lemmas.

LEMMA 5.8. *Let $P \Vdash_\rho^\sigma \Delta; \Gamma; \mathcal{G}$. $\text{co}(\Gamma)$ implies $\Delta = \emptyset$ or Δ contains only 1 or $\perp$ and $\sigma = \emptyset$.*

PROOF. See Appendix A.3.1. $\square$

LEMMA 5.9. *Let $P \Vdash_\rho^\emptyset \Delta; \Gamma; \mathcal{G}$ with $\Delta = \emptyset$ or Δ containing only 1 or $\perp$. We have that $\text{co}(\Gamma)$.*

PROOF. See Appendix A.3.2. $\square$

THEOREM 5.10. *Let $P \Vdash_{\rho}^{\emptyset} \Delta; \mathcal{G}$ and $\Delta = \emptyset$ or Δ contains only 1 or $\perp$. Then we have: (1) $P \rightarrow^* 0$; and (2) $\text{fuse}(\mathcal{G})$ is well formed.*

PROOF. By Propositions 2.3 and 5.5 and Theorem 4.9, together with Lemmas 5.8 and 5.9. $\square$

From Proposition 3.6 and Theorem 5.10, it follows that:

COROLLARY 5.11. *If $P \Vdash_{\rho}^{\emptyset} \Delta; \Gamma$ and $\Delta = \emptyset$ or Δ contains only 1 or $\perp$, then $\rho(P) \vdash_{\text{MP}} \Gamma$ and $\rho(P)$ is deadlock-free.*

Recall that Proposition 2.7 states that the MP typing discipline does not guarantee deadlock-freedom. Theorem 5.10 shows that the translation from CLL automatically identifies a set of deadlock-free MP processes.

6 HIGHER-ORDER CHANNEL PASSING

In this section, we lift the restrictions put in place in Section 3 on the $\otimes$ and $\mathcal{V}$ connectives, enabling CLL processes to perform full higher-order channel passing that can be mapped to MP processes with delegation. We follow the lines of Section 3 and Section 5, extending the framework and earlier results to this more general setting, emphasising crucial differences.

Channel Mappings. Full higher-order channel passing creates interleaved *multiple* sessions and instantiations of channels into input bound variables. For these reasons, we revise our mapping of Section 3, allowing for processes that send channels to hold multiple roles in the same multiparty session. We also account for bound names, where (bound) CLL channels are mapped to distinct MP session channels.

We present our mapping with two definitions: a mapping for cut-free processes (Definition 6.1) and a well-formedness condition for delegation (Definition 6.2). In the former, the mapping is identical to that of Section 3 but the MP channel identifier need not be unique among all channels. We also need to treat consistency of bound names. In the latter, we enforce that when the typing rule for delegation (i.e., the $\otimes$ rule) is applied, channels used by the subprocesses must implement a different principal role.

Definition 6.1 (Channel Mapping). Let $P \vdash_{\text{CL}} \Delta$ without using the cut rule. We define a channel to role-indexed channel mapping of P as a pair of mappings (σ, η) such that: **(1)** for all distinct $x, y \in \text{fn}(P)$, $\sigma(x) = s[p][q]$ and $\sigma(y) = s'[p'] [q']$ where if $s = s'$ then $p = p'$ and $q \neq q'$; **(2)** for all distinct $x, y \in \text{bn}(P)$, $\eta(x) = s[p][q]$ and $\eta(y) = s'[p'] [q']$ where $s \neq s'$ and $s, s' \notin \sigma$; and **(3)** for all distinct $x, y \in \text{bv}(P)$, $\eta(x) = x[p]$ and $\eta(y) = y[p']$.

The restrictions in Definition 6.1 allow for different CLL channels to be mapped to different MP session channels. However, within a given MP session we enforce that the principal role implemented by the cut-free process must be the same (which identifies a single participant with a cut-free process). We also ensure that different channels mapped to the same MP session do not have the same destination role as in the previous mapping (Definition 3.1).

Crucially, the η component of Definition 6.1 tracks all instances of channel output **(2)** and input **(3)**, where sent channels are mapped to fresh (binary) MP session channels for which one of the endpoints is delegated. Dually, received channels are mapped to variables with a role assignment.

Definition 6.2 (Delegation). Let $P \vdash_{\text{CL}} \Delta$ without using the cut rule, and (σ, η) be a mapping viz. Definition 6.1. We say that (σ, η) is *well formed* if the number of distinct MP channels in the image of σ is minimal and for each use of $\otimes$ in typing P we have

$$\frac{P_1 \vdash_{\text{CL}} \Delta_1, y:A \quad P_2 \vdash_{\text{CL}} \Delta_2, x:B}{\bar{x}(y).(P_1 \mid P_2) \vdash_{\text{CL}} \Delta_1, \Delta_2, x:A \otimes B}.$$

$$\begin{aligned}
\varphi_P(\overline{x}(y).(P_1 \mid P_2)) &\triangleq \begin{cases} (\nu c_\eta(y))c_\sigma(x)[p_\sigma(x)][d_\sigma(x)]\langle c_\eta(y)[d_\eta(y)] \rangle.(\varphi_P(P_1) \mid \varphi_P(P_2)) & \text{if } x \in \text{fn}(P) \\ (\nu c_\eta(y))x[d_\eta(x)]\langle c_\eta(y)[d_\eta(y)] \rangle.(\varphi_P(P_1) \mid \varphi_P(P_2)) & \text{if } x \in \text{bv}(P) \end{cases} \\
\varphi_P(x(y).Q) &\triangleq \begin{cases} c_\sigma(x)[p_\sigma(x)][d_\sigma(x)](y).\varphi_P(Q) & \text{if } x \in \text{fn}(P) \\ x[d_\eta(x)](y).\varphi_P(Q) & \text{if } x \in \text{bv}(P) \end{cases}
\end{aligned}$$

Fig. 6. Process mapping.

(Role Disjointness) $\forall z \in \Delta_1, z' \in \Delta_2, x: c_\sigma(z) = c_\sigma(z')$ implies $p_\sigma(z) \neq p_\sigma(z')$;

(Role/Destination Disjointness) $\forall z \in \Delta_1, z' \in \Delta, z \neq z'$. and $c_\sigma(z) = c_\sigma(z')$ imply $d_\sigma(z) \neq p_\sigma(z')$.

$P \vdash_{\text{CL}}^{(\sigma, \eta)} \Delta$ denotes (σ, η) is well formed for delegation w.r.t. $P \vdash_{\text{CL}} \Delta$.

The conditions of Definition 6.2 ensure that channels that are used within parallel compositions (due to delegation in CLL) are assigned different principal roles within the same MP session (the condition **(Role Disjointness)**). This ensures typability in MP. Moreover, the **(Role/Destination Disjointness)** condition forbids all principal roles used in the process from being used as destination roles within the same process (i.e., disallowing two endpoints of a communication within the same thread).

Now we define a main mapping from CLL to MP.

Definition 6.3 (Process Mapping). Given $P \vdash_{\text{CL}} \Delta$ without using the cut rule and a mapping (σ, η) according to Definition 6.1, we define the mapping from P to the MP process $\varphi_P(P)$, where $\varphi = \sigma \circ \eta$ according to the rules of Figure 6. We often omit the subscript P when clear from context.

The following example illustrates how our mapping for higher-order channel passing transforms CLL into MP processes and the conditions imposed on the mapping.

Example 6.4. Let $P \triangleq \overline{x}(y).(P_1 \mid P_2)$ with $P_1 \triangleq \overline{w_1}(1) \mid \overline{y}(2)$ and $P_2 \triangleq \overline{w_2}(3) \mid \overline{x}(5)$. Then assume (we write φ for $\sigma \circ \eta$):

$$\varphi(P) = (\nu s')s[p][q]\langle s'[q'] \rangle.(s_1[p_1][q_1]\langle 1 \rangle \mid s'[r][q']\langle 2 \rangle \mid s_2[p_2][q_2]\langle 3 \rangle \mid s[p][q]\langle 5 \rangle)$$

for some well-formed η and σ . Then $s_2 \neq s$ and by Definition 6.2:

- if $s = s_1$, $p \neq p_1$ (by (1)) and $q \neq p_1$ and $p \neq q_1$ (by (2)); and
- if $s_1 = s_2$, $p_1 \neq p_2$ (by (1)) and $q_1 \neq p_2$ and $q_2 \neq p_1$ (by (2)).

Note that a valid mapping does not exclude the cases where two destination roles can be same, i.e., $q = q_1$ and $q_1 = q_2$ are allowed when $s = s_1$ and $s_1 = s_2$. However, composition with such mappings will be subsequently excluded by the parallel composition rule in Figure 7.

Local Type Generation. Since cut-free processes may denote multiple roles, we generate a local typing *context* from each process, assigning local types to each role indexed channel. To generate the local type for a session output of the form $\overline{x}(y).(P \mid Q)$, where y occurs free in P but not in Q (and symmetrically for x), we produce the delegated session type from the usage of y in P . The two continuation processes are then inspected inductively. Since local type generation produces a local typing context instead of a single local type, we provide a combination operation $s[p]:T \diamond \Gamma$ that acts as a simple union, but for type assignments for $s[p]:T'$ in Γ results in an assignment in which T precedes T' . Note that, since P and Q use disjoint sets of channels (which are mapped to disjoint roles), applying $\diamond$ to the corresponding generated local types amounts to a simple set union.

$$\begin{array}{c}
\text{(comp)}_d \frac{P \vdash_{\text{CL}}^{\sigma, \eta} \Delta, x:A \quad Q \models_{\rho}^{\sigma', \eta'} \Delta', x:A^\perp; \Gamma; \mathcal{G} \quad (\dagger)}{(\nu x)(P \mid Q) \models_{\rho'}^{(\sigma' \cup \sigma) \setminus \{x\}, \eta \cup \eta'} \Delta, \Delta'; \Gamma, \llbracket P \rrbracket_\sigma^\eta; \mathcal{G} \cup \{P\}_\sigma^\eta} \\
\hline
(\dagger) \quad \begin{array}{l}
\text{(a) bound channel: } \rho' = \rho \cup (x, \mathbf{c}_\sigma(x)[p_\sigma(x)][d_\sigma(x)]) \\
\text{(b) role/destination match: } \mathbf{c}_\sigma(x) = \mathbf{c}_{\sigma'}(x) \wedge p_\sigma(x) = d_{\sigma'}(x) \wedge d_\sigma(x) = p_{\sigma'}(x) \\
\forall z \in \Delta, y \in \Delta'. \mathbf{c}_\sigma(z) = \mathbf{c}_{\sigma'}(y) \Rightarrow \\
\quad \text{(c) unique destination: } (d_\sigma(z) \neq d_{\sigma'}(y) \wedge d_\sigma(z), d_{\sigma'}(y) \notin \rho \wedge \\
\quad \text{(d) role/destination disjointness: } p_\sigma(z) \neq d_{\sigma'}(y) \wedge d_\sigma(z) \neq p_{\sigma'}(y) \wedge p_\sigma(z) \neq p_{\sigma'}(y)) \\
\text{(e) bound role/destination match: } \forall x \in \eta. \forall y \in \eta'. x = y \Rightarrow d_\eta(x) = p_{\eta'}(y) \vee p_\eta(x) = d_{\eta'}(y) \\
\text{(f) bound channel disjointness: } \forall x \in \eta. \forall y \in \varphi = \{\sigma', \eta', \rho\}. \mathbf{c}_\eta(x) \neq \mathbf{c}_\varphi(y)
\end{array}
\end{array}$$

Fig. 7. Parallel composition mapping with channel passing.

The operation $s[p]:T \diamond \Gamma$, that essentially appends a local type assignment $s[p]:T$ to those in Γ that match the session channel and role, is defined as follows (recall $\dagger$ denotes either $\{\uparrow, \downarrow, \oplus, \&\}$):

$$\begin{array}{ll}
s[p]:q\dagger(T) \diamond (\Gamma', s[p]:T') & \triangleq \Gamma', s[p]:q\dagger(T); T' \\
s[p]:q\dagger(T) \diamond \Gamma' & \triangleq s[p]:q\dagger(T); \text{end}, \Gamma' \quad \text{with } s[p] \notin \Gamma' \\
s[p]:\dagger q\{l_j : T_j\}_{j \in J} \diamond (\Gamma', s[p]:T') & \triangleq \Gamma', s[p]:\dagger q\{l_j : T_j\}_{j \in J} \\
s[p]:\dagger q\{l_j : T_j\} \diamond \Gamma' & \triangleq s[p]:\dagger q\{l_j : T_j\}, \Gamma' \quad \text{with } s[p] \notin \Gamma'.
\end{array}$$

Intuitively, $s[p]:T \diamond \Gamma$ is $\Gamma, s[p]:T$ if $s[p] \notin \Gamma$. Otherwise, we have that $\Gamma = \Gamma', s[p] : T'$, and thus we modify T' to $T; T'$ if T is an input or output; or, change T' to T if T is a selection or branching (since T will already contain all actions of role p by construction of the local type generation procedure).

Definition 6.5 (Local Type Generation). Let $P \vdash_{\text{CL}}^{(\sigma, \eta)} \Delta$, where all free and bound names are distinct. We generate a local typing context Γ such that $\eta(\sigma(P)) \vdash_{\text{MP}} \Gamma$ by induction on the structure of P , written $\llbracket P \rrbracket_\sigma^\eta$. Below we write $\mathbf{c}_\varphi(x)$, $\mathbf{p}_\varphi(x)$ and $\mathbf{d}_\varphi(x)$, where φ stands for σ if $x \in \text{fn}(P)$ and η otherwise; and c for $\mathbf{c}_\varphi(x)[p_\varphi(x)]$ if $x \in \text{fn}(P)$ and for x otherwise; q for $\mathbf{d}_\varphi(x)$; $\llbracket P \rrbracket_\sigma^\eta(y)$ denotes the binding for y in the generated context):

$$\begin{array}{ll}
\llbracket \mathbf{0} \rrbracket_\sigma^\eta & \triangleq \emptyset \\
\llbracket \bar{x}(y).(P \mid Q) \rrbracket_\sigma^\eta & \triangleq c:q\uparrow(\llbracket P \rrbracket_\sigma^\eta(y)) \diamond (\llbracket P \rrbracket_\sigma^\eta \diamond \llbracket Q \rrbracket_\sigma^\eta) \\
\llbracket x(y).P \rrbracket_\sigma^\eta & \triangleq c:q\downarrow(\llbracket P \rrbracket_\sigma^\eta(y)) \diamond (\llbracket P \rrbracket_\sigma^\eta \setminus \eta(y)) \\
\llbracket x.l_j; P \rrbracket_\sigma^\eta & \triangleq c:\oplus q\{l_j:\llbracket P \rrbracket_\sigma^\eta(c)\} \diamond \llbracket P \rrbracket_\sigma^\eta \\
\llbracket x.\text{case}\{l_i : P_i\}_{i \in I} \rrbracket_\sigma^\eta & \triangleq c:\& q\{l_i:\llbracket P_i \rrbracket_\sigma^\eta(c)\}_{i \in I} \diamond \llbracket P \rrbracket_\sigma^\eta,
\end{array}$$

where $\bar{T}$ denotes a dual type of T in a session with two roles p, q as follows:

$$\begin{array}{l}
\overline{p\uparrow(T);S} \triangleq q\downarrow(T);\bar{S} \quad \overline{p\downarrow(T);S} \triangleq q\uparrow(T);\bar{S} \\
\overline{\&p\{l_i : T_i\}_{i \in I}} \triangleq \oplus q\{l_i : \bar{T}_i\}_{i \in I} \quad \overline{\oplus p\{l_i : T_i\}_{i \in I}} \triangleq \&q\{l_i : \bar{T}_i\}_{i \in I} \quad \overline{\text{end}} \triangleq \text{end}.
\end{array}$$

We note that the carried type in outputs is dualised to match with session type duality.

Example 6.6. Consider

$$P \triangleq \bar{x}(y).(y(n).y(m).z\langle n \rangle.\mathbf{0} \mid \mathbf{0}) \quad Q \triangleq x(y).y\langle 0 \rangle.y\langle 1 \rangle.\mathbf{0}$$

with the following typings:

$$P \vdash_{\text{CL}} x:(\text{nat } \mathfrak{N} \text{ nat } \mathfrak{N} \perp) \otimes \mathbf{1}, z:\text{nat} \otimes \mathbf{1} \quad Q \vdash_{\text{CL}} x:(\text{nat} \otimes \text{nat} \otimes \mathbf{1}) \mathfrak{N} \perp.$$

We can produce mappings σ_1, η_1 and σ_2, η_2 such that (we write φ_i for $\sigma_i \circ \eta_i$):

$$\begin{aligned}\varphi_1(P) &= (\nu s')s[p][q]\langle s'[q'] \rangle; (s'[p']\langle q' \rangle(n); s'[p']\langle q' \rangle(m); s[t][v]\langle n \rangle; \mathbf{0} \mid \mathbf{0}) \\ \varphi_2(Q) &= s[q][p](y); y[p']\langle 0 \rangle; y[p']\langle 1 \rangle; \mathbf{0}\end{aligned}$$

and we have that

$$(\nu x)(P \mid Q) \models^{\sigma_3, \eta} z:\text{nat} \otimes 1; s[p]:T_1, s[q]:T_2, s[t]:T_3$$

with

$$\begin{aligned}T &= p' \uparrow (\text{nat}); p' \uparrow (\text{nat}); \text{end} \\ T_1 &= q \uparrow (T); \text{end} \\ T_2 &= p \downarrow (p' \uparrow (\text{nat}); p' \uparrow (\text{nat}); \text{end}); \text{end} \\ T_3 &= v \uparrow (\text{nat}); \text{end}.\end{aligned}$$

Note that mapping z to $s[p][t]$, for instance, would not allow for a valid composition of the two processes since we would have the role p of session s spread across two threads. Likewise, mapping z to $s[t][q]$ would disallow the composition of P and Q since it would require the two processes to share two distinct channel names.

Composition and Interconnection Networks. We define composition in tandem with the partial global types for delegation as other constructs are identical to Definition 5.1 (with a single session). We introduce the judgement $P \models^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$, where (σ, η) is a well-formed mapping according to Definition 6.2, following a similar pattern to the composition judgement of Section 3 and Section 5.

Definition 6.7 (Partial Global Types). Given $P \vdash_{\text{CL}}^{(\sigma, \eta)} \Delta$, we generate its partial global type w.r.t. s , written $\langle P \rangle_\sigma^\eta(s)$ by induction on the structure of P ($\#$ denotes $x \in \sigma \wedge c_\sigma(x) = s$):

$$\begin{aligned}\langle x\langle y \rangle. (Q_1 \mid Q_2) \rangle_\sigma^\eta(s) &\triangleq \begin{cases} p_\sigma(x) \rightsquigarrow d_\sigma(x) : \uparrow (\overline{\langle Q_1 \rangle_\sigma^\eta(y)}) . \text{fuse}(\langle Q_1 \rangle_\sigma^\eta(s), \langle Q_2 \rangle_\sigma^\eta(s)) & (\#) \\ \text{fuse}(\langle Q_1 \rangle_\sigma^\eta(s), \langle Q_2 \rangle_\sigma^\eta(s)) & \text{otherwise} \end{cases} \\ \langle x(y). Q \rangle_\sigma^\eta(s) &\triangleq \begin{cases} d_\sigma(x) \rightsquigarrow p_\sigma(x) : \downarrow (\langle Q \rangle_\sigma^\eta(y)) . \langle Q \rangle_\sigma^\eta(s) & (\#) \\ \langle Q \rangle_\sigma^\eta(s) & \text{otherwise.} \end{cases}\end{aligned}$$

Let C be the set of session channels in the image of σ . We denote by $\langle P \rangle_\sigma^\eta$ the set $\bigcup (\langle P \rangle_\sigma^\eta(s))_{s \in C}$.

We define $P \models^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$, where (σ, η) is a well-formed mapping, with the rule (changes w.r.t. $(\star)$ appear in **red**) in Figure 7. Following the Barendregt convention, we assume that if $x \in \text{bn}(P) \cap \text{bv}(Q)$ or $x \in \text{bv}(P) \cap \text{bn}(Q)$, then P and Q eventually exchange x . Conditions **(a)**–**(c)** are identical to (comp) in Figure 4 but where we refer explicitly to the mapping of CLL channels to MP session channels. Condition **(d)** ensures that P and Q are only connected via the channel x , as required by CLL, and that no roles are split across the two processes. Conditions **(e)** and **(f)** ensure that channels that are to be exchanged are mapped to fresh channels and with consistent role assignments in MP.

Example 6.8. The following processes are untypable in the global progress type system in Coppo et al. [18]:

$$\begin{aligned}P &\triangleq x\langle y \rangle. y(n). \mathbf{0} \vdash_{\text{CL}} x:(\text{int} \wp \perp) \otimes \mathbf{1} \\ Q &\triangleq z\langle w \rangle. w(1). \mathbf{0} \vdash_{\text{CL}} z:(\text{int} \otimes \mathbf{1}) \otimes \mathbf{1} \\ R &\triangleq x(y). z(w). w(n). y\langle n \rangle. \mathbf{0} \vdash_{\text{CL}} x:(\text{int} \otimes \mathbf{1}) \wp \perp, z:(\text{int} \wp \perp) \wp \perp.\end{aligned}$$

Given mappings $\sigma_1, \sigma_2, \sigma_3, \eta_1, \eta_2, \eta_3$, we obtain the following (with $\varphi_i = \sigma_i \circ \eta_i$):

$$\begin{aligned}\varphi_1(P) &= (\nu t)s[p][r]\langle t[t]; t[s][t](n); \mathbf{0} \\ \varphi_2(Q) &= (\nu t')s'[q][b]\langle t'[s]; t'[r][s]\langle 1 \rangle; \mathbf{0} \\ \varphi_3(R) &= s[r][p](y); s'[b][q](w); w[s](n); y[t]\langle n \rangle; \mathbf{0}\end{aligned}$$

$$\llbracket P \rrbracket_{\eta_1}^{\sigma_1} = s[p]:r \uparrow (s \uparrow (\text{int})) \quad \llbracket Q \rrbracket_{\eta_2}^{\sigma_2} = s'[q]:b \uparrow (r \downarrow (\text{int})) \text{ and } \llbracket R \rrbracket_{\eta_3}^{\sigma_3} = s[r]:p \downarrow (s \uparrow (\text{int})), s'[b]:q \downarrow (r \downarrow (\text{int})).$$

We showcase a form of systems that cannot be directly represented in the MCP system in Reference [11] as a single session.

Example 6.9 (Comparison with Carbone et al. [11, 14]). Consider the following CLL processes, where P_1 employs channel passing to send a session of type $\text{int} \otimes \mathbf{1}$ (we write φ_i for $\sigma_i \circ \eta_i$):

$$\begin{aligned}P_1 &\triangleq \bar{a}\langle y \rangle.(b(x).y\langle x \rangle \mid a\langle \rangle) \vdash_{\text{CL}} a:(\text{int} \otimes \mathbf{1}) \otimes \text{unit} \otimes \mathbf{1}, b:\text{int} \wp \perp \\ P_2 &\triangleq a(y).y\langle x \rangle.a\langle \rangle \vdash_{\text{CL}} a:(\text{int} \wp \perp) \wp \text{unit} \wp \perp, c:\mathbf{1} \\ P_3 &\triangleq b\langle 33 \rangle \vdash_{\text{CL}} b:\text{int} \otimes \mathbf{1}.\end{aligned}$$

We can define mappings $\sigma_1, \sigma_2, \sigma_3, \eta_1, \eta_2, \eta_3$ such that

$$\begin{aligned}\varphi_1(P_1) &= (\nu s')s[p][q]\langle s'[q_2]; (s[t][r](x); s'[p_2][q_2]\langle x \rangle \mid s[p][q]\langle \rangle) \\ \varphi_2(P_2) &= s[q][p](y); y[p_2](x); s[q][p]\langle \rangle \\ \varphi_3(P_3) &= s[r][t]\langle 33 \rangle.\end{aligned}$$

We can compose the three processes using our (comp_d) rule such that the corresponding global type G is $p \rightarrow q : (T).r \rightarrow t : (\text{int}).p \rightarrow q : (\text{unit})$. Note that MCP [11] cannot type this composition using G , since P_1 has actions of both p and t of s . To type a composition of this form in MCP, we have two options: (1) we force the actions on CLL channel b correspond to a separate MP session, thus requiring two global types ($G_1 = p \rightarrow q : (T).p \rightarrow q : (\text{unit})$ and $G_2 = r \rightarrow t : (\text{int})$) to type the corresponding processes, or (2) we separate role t into an independent fourth process $P_4 = s[t][r](x)$, removing the communication from P_1 and requiring an additional *thread* at the start of the session (note the dependency between the input from r and the output on s' is lost).

Consistency and Acyclicity. Given that a cut-free process might contain many sessions and roles per session (i.e., multiple threads), we extend the notion of a thread and thread-preservation.

Informally, we can regard a cut-free process P as a thread if (1) each sequential subterm of P contains only one role per session; and (2) in each delegation subterm of P of the form $\bar{x}\langle y \rangle.(Q \mid R)$, the delegated channel y is allocated to a different session. Condition (1) means that the mapping builds the longest (i.e., with the most communication steps) possible typable session, and condition (2) avoids self-delegation. This is consistent with the conditions of Definition 6.1 and Definition 6.2.

To state a similar property to Theorem 3.9, we move from a single session to multiple sessions, where bijective renaming φ ensures that (1) if different channels map to the same session, their destination roles differ; and, (2) if the same session channel appears in two parallel prefixes, their principal and destination roles are pairwise distinct. Since judgement $P \models_{\rho}^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$ relies on the channel mappings in Definitions 6.1 and 6.2, we prove the following property that corresponds to Theorem 3.9 in this channel-passing setting, under the assumption that ρ, σ and η conform the channel mappings.

THEOREM 6.10 (CONSISTENCY). *thedeuniqueness Assume $P \vdash_{\text{CL}} \Delta$. If $\varphi(P)$ is typable by Γ , i.e., $\varphi(P) \vdash_{\text{MP}} \Gamma$ and ρ, σ and η satisfy the conditions in Definitions 6.1 and 6.2 and $\varphi = \sigma \circ \eta \circ \rho$, then $P \models_{\rho}^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$.*

PROOF. See Appendix A.4.1. □

The property above, while weaker than the uniqueness property of Theorem 3.9, provides a form of consistency between the channel mappings and the (comp_d) rule in Figure 7.

PROPOSITION 6.11. *Let $P \models_{\rho}^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$ and $\Delta = \emptyset$ or Δ contains only 1 or $\perp$. There exists a single well-formed global type G , such that $G = \text{fuse}_s(\mathcal{G})$ where $\text{fuse}_s(\mathcal{G})$ denotes fusion of all partial global types for session s in $\mathcal{G}$.*

PROOF. Identical to Proposition 5.5 due to the fact that instances of session delegation that are composed in P are mapped to independent and complete global types. $\square$

We prove that delegation does not add to the connection graph since delegation denotes a distinct multiparty session (i.e., the analogue of Theorem 6.12).

THEOREM 6.12. *Let $P \models_{\rho}^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$ and $\Delta = \emptyset$ or Δ containing only 1 or $\perp$. The interconnection network graph for $\text{fuse}_s(\mathcal{G})$ for each session s is acyclic.*

PROOF. Assume to the contrary that the connection graph for G has a cycle. We have two kinds of cycles: a sequence of edges of the form either:

- (1) *Triangle:* $(p, q), (p, r), (q, r)$; or
- (2) *Diamond:* $(p, q), (p, r), (q, t_1), \dots, (t_n, s), (r, v_1), \dots, (v_m, s)$.

We show that Δ cannot be empty or contain only 1 or $\perp$ in either case, deriving a contradiction. The proof follows the general lines of that of Theorem 5.7, but with a more involved case analysis since a cut-free process may now implement multiple roles (which by the restrictions of Definitions 6.1 and 6.2 will not be connected), and thus we need to account for all possibilities of roles being implemented by the same process. In particular, the case for “diamond” shaped connections (**Case (2)** below) of the form $(p, q), (p, r), (q, t_1), \dots, (t_n, s), (r, v_1), \dots, (v_m, s)$ requires us to also account for the fact that r and q might be implemented by the same cut-free process, and similarly for q and v_1 , r and t_1 , and so on.

Case (1): Assume the connection graph for G contains a triangle. Since (p, q) is in the connection graph, we know that roles p and q cannot be implemented in the same process thread. Similarly for p and r and q and r . Thus, we must have at least three process threads (P_1, P_2 , and P_3), one for each role. Without loss of generality, assume $P_1 \vdash_{\text{CL}}^{\sigma_0} x:A, y:B$ with $\sigma_0(x) = s[p][q]$ and $\sigma(y) = s[p][r]$. $P_2 \vdash_{\text{CL}}^{\sigma_1} x:A^{\perp}, z:C$ with $\sigma_1(x) = s[q][p]$ and $\sigma_1(z) = s[q][r]$. It is then immediate that we cannot find any P_3 implementing r (to fully empty the context) since it would have to share two channel names with the composition of P_1 and P_2 , which is not a well-formed composition according to $(\dagger)$ in rule (comp_d) .

Case (2): Assume the connection graph for G contains a diamond. We already know by Theorem 5.7 that when all roles are implemented by separate process threads that we cannot form a diamond. Hence, by Definitions 6.1 and 6.2, it must be the case that unconnected roles in the graph are implemented in the same process.

We note that if q and r are implemented by the same process thread, then we cannot find a closing instance of p (since we would need to compose two processes sharing two channel names, which is disallowed by rule (comp_d) , or with ill-formed mappings according to $(\dagger)$ in the same rule).

We proceed by case analysis on (n, m) .

Case (2-1) $n = 0$ and $m = 0$: We cannot find a closed instance of the network since either it is the case that each role is assigned to each thread (and so Theorem 5.7 applies) or p and s are implemented by the same thread. If this were the case, then we must have $P_1 \vdash_{\text{CL}}^{\sigma_0} x:A, y:B, z:C, w:D$

such that $\sigma_0(x) = s[p][q]$, $\sigma_0(y) = s[p][r]$, $\sigma_0(z) = s[s][r]$ and $\sigma_0(w) = s[s][q]$. This is not possible by Definitions 6.1 and 6.2.

Case (2-2) $n = 0$ and $m = m' + 1$: We have established that (p and q), (p and r), and (q and r) cannot be implemented by the same thread. If q and v_1 were implemented by the same thread, then we cannot find a closed instance of this network since we need to compose with the (distinct) implementations of p and r which themselves are connected (by the assumption of the existence of a diamond connection). Thus we need to compose a process mapped to $(s[q][p], s[v_1][r], \text{and } s[v_1][v_2])$ with a process mapped to $(s[p][q] \text{ and } s[p][r])$ and another mapped to $(s[r][p] \text{ and } s[r][v_1])$. If we compose the first with the second, then we cannot compose with the third since two channel/role assignment pairs $(s[r][p] \text{ and } s[r][v_1])$ are shared, which is forbidden by $(\dagger)$ in rule (comp_d) . A similar reasoning applies to the other ways of composing the processes. The same reasoning applies if q and v_i were implemented by the same thread.

Case (2-3) $n = n' + 1$ and $m = m' + 1$: We begin with the case where q and v_1 are implemented by the same thread and r and t_1 are implemented by the same thread, which are obviously uncomposable according to rule (comp_d) , as two distinct channels are shared. The same reasoning applies as we increase n and m . If t_i and v_i are implemented by the same thread, then we must have processes sharing two distinct channels, which is impossible. If t_n and v_m are mapped to the same process, then the implementation of s must share two channels with this process, which is also a contradiction. Finally, if p and s are the same process, the same reasoning described in **Case (2-1)** case applies. $\square$

LEMMA 6.13. *Let $P \models_{\rho}^{\emptyset, \eta} \Delta; \Gamma; \mathcal{G}$. $\text{co}(\Gamma)$ implies $\Delta = \emptyset$ or Δ containing only 1 or $\perp$.*

PROOF. We note that the renamings ensure that the two endpoints of an interaction cannot be implemented by the same single-thread process and that bound-names involved in delegation denote linear interactions along different session channels. Hence a similar argument using the contradiction is applicable (with a case analysis for delegations) as the proof of Lemma 5.8. See Appendix A.4.2 for the details. $\square$

LEMMA 6.14. *Let $P \models_{\rho}^{\emptyset, \eta} \Delta; \Gamma; \mathcal{G}$ with $\Delta = \emptyset$ or Δ containing only 1 or $\perp$. We have that $\text{co}(\Gamma)$.*

PROOF. We account for the additional case of delegation by noting that the sent channel has a dual behaviour to the received channel (and we generate compatible endpoint types in Γ). The result then follows using a similar argument as the proof of Lemma 5.9. $\square$

THEOREM 6.15. *Let $P \models_{\rho}^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$ and $\Delta = \emptyset$ or Δ containing only 1 or $\perp$. Then we have: (1) $P \rightarrow^* 0$; and (2) $\text{fuse}_s(\mathcal{G})$ at each session s is well formed and deadlock-free.*

PROOF. (1) follows from [9, 60] and (2) follows from Proposition 6.11. $\square$

COROLLARY 6.16. *If $P \models_{\rho}^{\emptyset, \eta} \Delta; \Gamma; \mathcal{G}$ and $\Delta = \emptyset$ or Δ contains only 1 or $\perp$, then $\eta(\rho(P))$ is deadlock-free.*

7 REPLICATION

We account for CLL replication within our framework of MP types. Typically, MP processes (and types) do not explicitly account for replication, thus instead of extending the MP process calculus with a non-standard form of MP replication that corresponds to the binary sessions of CLL, we reason at the level of types. From the point of view of interconnectability *within* a linear session, the replication construct consists of the ability to “repeat” existing interconnections in *fresh*, independent sessions. While a replicated session may be used by *many* processes (due to the (cut^1) rule

that allows for many threads to use a replicated name, akin to a multicut for non-linear sessions), each session replica is necessarily independent since rule (!) does not allow a replicated session to rely on other linear sessions. Hence the addition of replication does not change the nature of interconnectability of linear sessions in CLL—each client of a replicated session is connected to a distinct, isolated instance. Given that we focus on the interconnectability of CLL (hence generation of global and local types), we establish our main results by extending global and local *types* (but not the MP process syntax) with a form of replication.

Channel Mapping. We define our mapping as in Definition 6.1, but we allow for different session names in the same thread and multiple destination roles for replicated channel names (otherwise the mapping of replicated inputs would be degenerate). Given $P \vdash_{\text{CL}} \Xi; \Delta$ we enforce that the channels in Ξ and Δ be mapped to distinct MP session channels. Moreover, we require that all linear channels that deal with replicated sessions be mapped to distinct MP sessions. We write $\text{posBang}(A)$ iff $!/?$ occurs to the right of $\otimes$ or $\wp$, or in a selection or branching in A . This predicate ensures that the behaviour specified by A aims to offer or use a replicated session but not delegate such a session. For the case of replicated inputs, p_σ denotes a single role and d_σ denotes a set of roles. The rest is unchanged.

Definition 7.1 (Channel Mapping—Replication). Let $P \vdash_{\text{CL}} \Xi; \Delta$ without using the cut rule. We define a channel to role-indexed channel mapping of P as a pair of mappings (σ, η) such that: (1) for all distinct $x, y \in \text{fn}(P)$, then if $x \in \Xi$ and $y \in \Delta$ or $x, y \in \Xi$, then $c_\sigma(x) \neq c_\sigma(y)$. If $x, y \in \Delta$, then $\sigma(x) = s[p][q]$ and $\sigma(y) = s'[p'][q']$, where if $s = s'$, then $p = p'$ and $q \neq q'$; (2) for all distinct $x, y \in \text{bn}(P)$, $\eta(x) = s[p][q]$ and $\eta(y) = s'[p'][q']$, where $s \neq s'$ and $s, s' \notin \sigma$; (3) for all distinct $x, y \in \text{bv}(P)$, $\eta(x) = x[p]$ and $\eta(y) = y[p']$; and (4) $\forall x:A, y:B \in \Delta$ with $x \neq y$ if $\text{posBang}(A)$, then $c_\sigma(x) \neq c_\sigma(y)$.

Clause (1) above, beyond the identical condition from Definition 6.1, ensures that replicated CLL session channels are mapped to distinct MP sessions that do not clash with those for linear CLL channels; clauses (2) and (3) deal with the treatment of bound channel names and variables, as in Definition 6.1; finally, clause (4) ensures that different CLL channels that aim to offer or use replicated session behaviours (i.e., whose types have occurrences of $!$ or $?$ that are not delegated) are not mapped to clashing MP session channels.

Global Types and Type Generation. We extend the syntax of global types with the constructs for replication. Partial global types with a dedicated replication construct are $\tilde{p} \rightsquigarrow q : ! (T).G$, with the corresponding dual $p \rightsquigarrow q : ? (T).G$ which fuse to the corresponding $\tilde{p} \rightarrow q : * (T).G$ global type, denoting that role q hosts the replicated behaviour T , to be used by roles $\tilde{p}$ an arbitrary (but finite) number of times.

We then extend the partial global type generation to account for replication as follows.

Definition 7.2 (Type Generation). Let $P \vdash_{\text{CL}}^{(\sigma, \eta)} \Xi; \Delta$, with all free and bound names distinct. We generate a set of role-indexed channels and types and partial global types w.r.t. a multiparty session channel s , $\llbracket P \rrbracket_\sigma^\eta$ and $\langle P \rangle_\sigma^\eta(s)$, respectively, by induction on the structure of P , as follows (we assume the same notations of Definition 6.5 and 6.7):

$$\begin{aligned} \llbracket !x(y).P \rrbracket_\sigma^\eta &\triangleq c:d_\varphi(x)!(\llbracket P \rrbracket_\sigma^\eta(y)) \uplus (\llbracket P \rrbracket_\sigma^\eta \setminus \eta(y)) \\ \llbracket \bar{x}(y).P \rrbracket_\sigma^\eta &\triangleq c:d_\varphi(x)?(\llbracket P \rrbracket_\sigma^\eta(y)) \uplus (\llbracket P \rrbracket_\sigma^\eta \setminus \eta(y)) \\ \langle !x(y).Q \rangle_\sigma^\eta(s) &\triangleq \begin{cases} d_\sigma(x) \rightsquigarrow p_\sigma(x) : !(\langle Q \rangle_\sigma^\eta(y)).\langle Q \rangle_\sigma^\eta(s) & (x \in \sigma \wedge c_\sigma(x) = s) \\ \langle Q \rangle_\sigma^\eta(s) & \text{otherwise} \end{cases} \\ \langle \bar{x}(y).Q \rangle_\sigma^\eta(s) &\triangleq \begin{cases} p_\sigma(x) \rightsquigarrow d_\sigma(x) : ?(\langle Q \rangle_\sigma^\eta(y)).\langle Q \rangle_\sigma^\eta(s) & (x \in \sigma \wedge c_\sigma(x) = s) \\ \langle Q \rangle_\sigma^\eta(s) & \text{otherwise.} \end{cases} \end{aligned}$$

We can then define composition $\Vdash_{\rho}^{\sigma, \eta} \Xi; \Delta; \Gamma; \mathcal{G}$ as in Figure 7 extending the conditions (a)–(f) considering d_{σ} and d_{η} as sets (and the appropriate checks for set membership and non-membership).

LEMMA 7.3 (PARTIAL GLOBAL TYPES). *A grammar of partial global types generated under Definitions 7.1 and 7.2 for each session is given as*

$$\begin{aligned} G ::= & \text{end} \mid \tilde{p} \rightsquigarrow q : !(T).\text{end} \mid p_1 \rightsquigarrow q : ?(T).p_2 \rightsquigarrow q : ?(T) \dots p_n \rightsquigarrow q : ?(T).\text{end} \\ & \mid p \rightsquigarrow q : \uparrow(T).G \mid p \rightsquigarrow q : \downarrow(T).G \mid p \rightsquigarrow q : \oplus\{l_j : G_j\}_{j \in J} \mid p \rightsquigarrow q : \&\{l_j : G_j\}_{j \in J}. \end{aligned}$$

PROOF. By Definition 7.1(1,3,4), if x in $!x(z).P$ is assigned to s , P does not contain any channels mapped to s . Hence if typable Q contains both $C_1[!x(z).P]$ (with type of x is $!$) and $C_2[\bar{x}\langle y \rangle.R]$ (with type of x is $?$) as its subterms, (1) C_i does not contain replication mapped to s ; (2) C_1 does not contain channels of $?$ type mapped to s ; and (3) we can set C_2 so that it not contain channels of $?$ type mapped to s . Thus at MP session s , $\tilde{p} \rightsquigarrow q : !(T).\text{end}$ is generated from $!x(z).P$, and $p_1 \rightsquigarrow q : ?(T).p_2 \rightsquigarrow q : ?(T) \dots p_n \rightsquigarrow q : ?(T).\text{end}$ with $p_j \in \tilde{p}$ is generated from $\bar{x}\langle y \rangle.R$ where the condition $p_j \in \tilde{p}$ is ensured by the conditions of $\Vdash_{\rho}^{\sigma, \eta}$. $\square$

Fusing and Complete Global Types. We can now define fuse over partial global types given in Lemma 7.3. The definition is extended as follows (with $T_1 \leq T_2$, $p \in \tilde{p}$ and omitting the congruence cases where the the fuse operation pushes under both the partial and complete replication prefixes):

$$\begin{aligned} \text{fuse}(\tilde{p} \rightsquigarrow q : !(T).\text{end}, \text{end}) &= \tilde{p} \rightarrow q : *(T).\text{end} \\ \text{fuse}(\tilde{p} \rightarrow q : *(T).\text{end}, \text{end}) &= \tilde{p} \rightarrow q : *(T).\text{end} \\ \text{fuse}(\tilde{p} \rightsquigarrow q : !(T_1).G_1, p \rightsquigarrow q : ?(T_2).G_2) &= \text{fuse}(\tilde{p} \rightsquigarrow q : !(T_1).G_1, G_2) \quad p \in \tilde{p}. \end{aligned}$$

LEMMA 7.4 (COMPLETE GLOBAL TYPES). *A grammar of complete global types generated under Definitions 7.1 and 7.2 with fuse for each session is given as*

$$G ::= \text{end} \mid \tilde{p} \rightarrow q : *(T).\text{end} \mid p \rightarrow q : (T).G \mid p \rightarrow q : \{l_j : G_j\}_{j \in J}.$$

PROOF. We only need to check that the three rules for replication defined above produce a global type of the form $\tilde{p} \rightarrow q : *(T).\text{end}$. By Lemma 7.3, the third rule is replaced by

$$\text{fuse}(\tilde{p} \rightsquigarrow q : !(T_1).\text{end}, p \rightsquigarrow q : ?(T_2).G_2) = \text{fuse}(\tilde{p} \rightsquigarrow q : !(T_1).\text{end}, G_2) \quad p \in \tilde{p},$$

where G_2 is either end or $p_1 \rightsquigarrow q : ?(T_2) \dots p_n \rightsquigarrow q : ?(T_2).\text{end}$, where $p_i \in \tilde{p}$. If $G_2 = \text{end}$, then the next matched rule is the first rule, which produces $\tilde{p} \rightarrow q : *(T).\text{end}$. We repeat the third rule until we reach $p_n \rightsquigarrow q : ?(T_2).\text{end}$ to then produce $\tilde{p} \rightarrow q : *(T).\text{end}$. $\square$

Note that the second rule of $\text{fuse}(G)$ would be used when the generated global type is in the form of the second line of the grammar in Lemma 7.3.

Example 7.5 (Racing on a Replicated Session). We illustrate the concepts pertaining to replication by considering the following processes (as in earlier examples, we assume basic value passing):

$$\begin{aligned} P &\triangleq !x(y).y(z_1).y(z_2).y\langle z_1 + z_2 \rangle \\ Q &\triangleq \bar{x}\langle y \rangle.y\langle 1 \rangle.y\langle 2 \rangle.y(z_3) \\ R &\triangleq \bar{x}\langle y \rangle.y\langle 2 \rangle.y\langle 1 \rangle.y(z_4), \end{aligned}$$

where $P \vdash_{\text{CL}} ; x : !(\text{Int} \wp \text{Int} \wp \text{Int} \otimes 1)$, $Q \vdash_{\text{CL}} ; x : ?(\text{Int} \otimes \text{Int} \otimes \perp)$, and $R \vdash_{\text{CL}} ; x : ?(\text{Int} \otimes \text{Int} \otimes \perp)$. We define mappings σ , σ_1 , σ_2 and η , η_1 , η_2 such that:

$$\begin{aligned} \llbracket P \rrbracket_{\sigma}^{\eta} &= s[p] : \{q, r\} ! (p_0 \uparrow (\text{Int}); p_0 \uparrow (\text{Int}); p_0 \downarrow (\text{Int})) \\ \llbracket Q \rrbracket_{\sigma_1}^{\eta_1} &= s[q] : p ? (p_0 \uparrow (\text{Int}); p_0 \uparrow (\text{Int}); p_0 \downarrow (\text{Int})) \\ \llbracket R \rrbracket_{\sigma_2}^{\eta_2} &= s[r] : p ? (p_0 \uparrow (\text{Int}); p_0 \uparrow (\text{Int}); p_0 \downarrow (\text{Int})). \end{aligned}$$

Generating the following global types:

$$\begin{aligned} \langle P \rangle_\sigma^\eta &= \{q, r\} \rightsquigarrow q : !(p_0 \uparrow(\text{Int}); p_0 \uparrow(\text{Int}); p_0 \downarrow(\text{Int})).\text{end} \\ \langle Q \rangle_{\sigma_1}^{\eta_1} &= \{q, r\} \rightsquigarrow q : ?(p_0 \uparrow(\text{Int}); p_0 \uparrow(\text{Int}); p_0 \downarrow(\text{Int})).\text{end} \\ \langle R \rangle_{\sigma_2}^{\eta_2} &= \{q, r\} \rightsquigarrow q : ?(p_0 \uparrow(\text{Int}); p_0 \uparrow(\text{Int}); p_0 \downarrow(\text{Int})).\text{end} \end{aligned}$$

It is then straightforward that

$$\text{fuse}(\langle P \rangle_\sigma^\eta, \text{fuse}(\langle Q \rangle_{\sigma_1}^{\eta_1}, \langle R \rangle_{\sigma_2}^{\eta_2})) = \{q, r\} \rightarrow p : *(p_0 \uparrow(\text{Int}); p_0 \uparrow(\text{Int}); p_0 \downarrow(\text{Int})).\text{end}.$$

Local Types and Liveness. To prove the main results, we now extend the syntax of local types with the constructs $\tilde{p}!(T)$ and $p?(T)$, denoting the type of a participant that expects to input from $\tilde{p}$, afterwards spawning a replica of type T , and the type of a participant that will use a replicated session offered by p (by sending it a fresh session) of type T .

Similarly, the syntax of binary types T extends with $!(T)$ and $?(T)$ and the duality of binary types is extended to $\overline{!(T)} = ?(\overline{T})$ and $\overline{?(T)} = !(\overline{T})$. Then the partial projection is defined as an input and an output of Definition 2.4, respectively.

The projection of $G = \tilde{s} \rightarrow r : *(T).\text{end}$ is defined as follows:

$$G \upharpoonright p = r?(T) \text{ if } p \in \tilde{s}; \quad G \upharpoonright p = \tilde{s}!(T) \text{ if } p = r; \quad G \upharpoonright p = \text{end} \text{ otherwise.}$$

The labels for LTSs are extended with $p\tilde{q}!(T)$ and $pq?(T)$, and the LTS rules of the local types are defined as

$$q?(T) \xrightarrow{pq?(T)} \text{end} \quad \tilde{q}!(T) \xrightarrow{p\tilde{q}!(T)} \tilde{q}!(T).$$

We extend the duality of labels as $\overline{p\tilde{q}!(T)} = q_i p?(T)$ and $\overline{q_i p?(T)} = p\tilde{q}!(T)$ with $q_i \in \tilde{q}$. Then the transitions of the configurations do not change. The semantics of G is defined as

$$(\text{expo}) \quad \tilde{p} \rightarrow q : *(T).\text{end} \xrightarrow{pq?(T) \cdot q\tilde{p}!(T)} \tilde{p} \rightarrow q : *(T).\text{end} \quad \text{with } p \in \tilde{p}.$$

We extend Definition 4.7 by adding the following cases:

4. if $\ell = pq?(T)$ there exists $C \xrightarrow{\vec{\ell}'} C' \xrightarrow{\ell} \vec{\ell} C''$;
5. if $\ell = p\tilde{q}!(T)$ for all $C \xrightarrow{\vec{\ell}'} C''$ such that $C'' = (T_p')_{p \in \mathcal{P}}, T_p' \xrightarrow{\ell} T_p''$.

Theorem 4.9 is updated replacing (DF) by the following liveness property, which allows for leftover replicated types (akin to Definition 2.2). Note that a difference from Theorem 4.9 is that $? \text{-output}$ is treated as an output and a selection since the replication is always available.

Definition 7.6 (Live). $C = (T_0)_p \in \mathcal{P}$ is *live* if for all $C \xrightarrow{\vec{\ell}} C_1 = (T_p)_{p \in \mathcal{P}}$, if $T_p \xrightarrow{\ell} T_p'$ and ℓ is not $!$, there exists $C' = (T_p'')_{p \in \mathcal{P}}$ such that $C_1 \xrightarrow{\vec{\ell}'} C'$ and (1) $C' \xrightarrow{\ell \cdot \vec{\ell}} C''$ if ℓ is an output or a selection or $? \text{-output}$; (2) $C' \xrightarrow{\vec{\ell} \cdot \ell} C''$ if ℓ is an input; or (3) $C' \xrightarrow{\vec{\ell}' \cdot \ell'} C''$ if $\ell = pq \triangleright l$ with some $\ell' = pq \triangleright l'$.

For the main theorem, we replace Definition 4.8 by Definition 7.6 to account for replication by mirroring the notion of live process.

THEOREM 7.7. Let $P \Vdash_\rho^{\emptyset, \eta} \cdot; \Delta; \Gamma; \mathcal{G}$ and $\Delta = \emptyset$ or Δ contains only 1 or $\perp$. We have: (1) If *live*(P), then $P \rightarrow P'$; (2) at each session s , $\text{fuse}_s(\mathcal{G})$ is well formed and live; and (3) the ING is acyclic.

PROOF. (1) follows from References [9, 60]; (2) we first note that replication is mapped to distinct multiparty sessions. Lemma 7.4 proves well-formedness. By Lemma 7.3 and Lemma 7.4, a complete replicated global type at each s is the form of $\tilde{p} \rightarrow q : *(T).\text{end}$. Then by the rules of LTSs defined

above, if $\ell = \text{pq?}(T)$ (?-output, which is the case of (1) in Definition 7.6), then $C' \xrightarrow{\ell.\bar{\ell}} C''$ since C contains $\text{pq!}(T)$ by the definition of projection. Hence for each s , $\text{fuse}_s(\mathcal{G})$ is live; and (3) follows the same argument as Theorem 5.7, noting that each replicated session corresponds to a different session. $\square$

The uniqueness theorem for replication, cf. Theorems 3.9 and 6.10, can be obtained under the condition where each replicated channel is assigned to a new session.

Example 7.8. The following processes (from an example of Carbone et al. [14]) are untypable in the global progress type system of [18].

$$\begin{aligned} P &\triangleq !x(y).y(n) \vdash_{\text{CL}} x:(\text{int } \wp 1) \quad Q \triangleq !z(w).w(1) \vdash_{\text{CL}} z:(\text{int } \otimes 1) \\ R &\triangleq \bar{x}(y).\bar{z}(w).w(n).y(n).0 \vdash_{\text{CL}} \cdot; x:(\text{int } \otimes \perp), z:(\text{int } \wp \perp). \end{aligned}$$

We define mappings $\sigma, \sigma_1, \sigma_2$ and η, η_1, η_2 such that

$$\langle P \rangle_{\sigma}^{\eta} = s[p]:r!(p_0 \uparrow (\text{int})) \quad \langle Q \rangle_{\sigma_1}^{\eta_1} = s'[q]:b!(r \downarrow (\text{int})) \quad \langle R \rangle_{\sigma_2}^{\eta_2} = s[r]:p?(p_0 \uparrow (\text{int})), s'[b]:q?(r \downarrow (\text{int})).$$

Then $(\nu x, z)(P \mid Q \mid R) \Vdash \cdot; \langle P \rangle_{\sigma}^{\eta}, \langle Q \rangle_{\sigma_1}^{\eta_1}, \langle R \rangle_{\sigma_2}^{\eta_2}$.

8 MULTICUT IN CLL

The inability to compose processes that interact by sharing more than one channel—often dubbed *multicut*—significantly limits the interconnection networks in CLL. Logically, such a form of unrestricted multicut is *unsound* and operationally results in deadlocks. Consider the following CLL processes:

$$P \triangleq y(x).z(7).0 \quad R \triangleq w(x).y\langle a \rangle.0 \quad Q_1 \triangleq z(x).w\langle \text{tt} \rangle.0 \quad Q_2 \triangleq w\langle \text{tt} \rangle.z(x).0.$$

We have that P is typed in a context $\Delta = y:\text{str } \wp \perp, z:\text{int } \otimes 1$, R is typed in a context $\Delta' = y:\text{str } \otimes 1, w:\text{bool } \wp \perp$ and both Q_1 and Q_2 are typed in a context $\Delta'' = z:\text{int } \wp \perp, w:\text{bool } \otimes 1$. We can observe that the composition $(\nu y, z, w)(P \mid Q_1 \mid R)$ is clearly deadlocked (viz. Section 2.2). However, the composition $(\nu y, z, w)(P \mid Q_2 \mid R)$ is safe, with the processes reducing in three steps to 0 . Intuitively, despite Q_1 and Q_2 both implementing the same sessions, their (identical) typings do not distinguish the sequential orderings of actions.

As discussed in Section 3, the framework of MP can distinguish such orderings. Thus we may appeal to the higher discriminating power of PMC to eliminate these unsafe (i.e., deadlocking) multicuts. For instance, the following MP processes can be mapped from the CLL processes above:

$$\begin{aligned} P' &\triangleq s[p][r](x); s[p][q]\langle 7 \rangle & R' &\triangleq s[r][q](x); s[r][p]\langle a \rangle \\ Q'_1 &\triangleq s[q][p](x); s[q][r]\langle \text{tt} \rangle & Q'_2 &\triangleq s[q][r]\langle \text{tt} \rangle; s[q][p](x). \end{aligned}$$

Then the (partial) global types generated from P' , R' , and Q'_1 are *not* PMC, whereas those from P' , R' , and Q'_2 are as follows:

$$\begin{aligned} \langle P' \rangle &\triangleq r \leadsto p: \downarrow(\text{str}).p \leadsto q: \uparrow(\text{int}).\text{end} & \langle Q'_1 \rangle &\triangleq p \leadsto q: \downarrow(\text{int}).q \leadsto r: \uparrow(\text{bool}).\text{end} \\ \langle R' \rangle &\triangleq q \leadsto r: \downarrow(\text{bool}).r \leadsto p: \uparrow(\text{str}).\text{end} & \langle Q'_2 \rangle &\triangleq q \leadsto r: \uparrow(\text{bool}).p \leadsto q: \downarrow(\text{int}).\text{end}. \end{aligned}$$

We thus make use of PMC to develop two deadlock-free multicut rules for CLL. We allow processes to share multiple channels holding dual types of each other (as the multicut rule in Reference [2]) but restrict composition by requiring the induced partial global types to be *fuseable* (or PMC), recovering deadlock-freedom.

Definition 8.1 (Multicut—No Name Passing or Replication). Using the mapping of Section 3, we re-define the judgement $P \Vdash_{\rho}^{\sigma} \Delta; \mathcal{G}$ to produce a multicut rule, defined in Figure 8 where “fuse($\mathcal{G}_1, \mathcal{G}_2$)

$$\begin{array}{c}
P \Vdash_{\rho}^{\sigma} \Delta, x_1:A_1, \dots, x_n:A_n; \mathcal{G}_1 \quad Q \Vdash_{\rho'}^{\sigma'} \Delta', x_1:A_1^{\perp}, \dots, x_n:A_n^{\perp}; \mathcal{G}_2 \\
\mathcal{G} = \text{fuse}(\mathcal{G}_1, \mathcal{G}_2) \text{ defined} \quad \rho' \cap \rho = \emptyset \quad (\ddagger) \\
\hline
(\text{MCut}) \quad \frac{(\nu x_1, \dots, x_n)(P \mid Q) \Vdash_{\rho''}^{(\sigma' \cup \sigma) \setminus \{x_1, \dots, x_n\}} \Delta, \Delta'; \mathcal{G}}{(\ddagger)}
\end{array}$$

$(\ddagger)$ (a) bound channels: $\rho'' = \rho \cup \rho' \cup_i (x_i, s[p_{\sigma}(x_i)][d_{\sigma}(x_i)])$
(b) role/destination match: $p_{\sigma}(x_i) = d_{\sigma'}(x_i) \wedge d_{\sigma}(x_i) = p_{\sigma'}(x_i)$
(c) unique destination: $\forall y \in \Delta. z \in \Delta'. (d_{\sigma}(y), d_{\sigma'}(z) \notin \rho, \rho' \wedge$
(d) role/destination disjointness: $p_{\sigma}(y) \neq p_{\sigma'}(z) \wedge p_{\sigma}(y) \neq d_{\sigma'}(z) \wedge p_{\sigma'}(z) \neq d_{\sigma}(y)$

Fig. 8. Multicut rule.

defined” means that for all $G_i \in \mathcal{G}_1 \cup \mathcal{G}_2$, $\text{fuse}(\dots (\text{fuse}(G_1, G_2), G_3)) \dots, G_n)$ is defined (Definition 4.10).

Rule (MCut) is symmetric, matching a generalised cut rule. We highlight the differences of (MCut) from (comp) and (comp_d) with **red letters**. Clauses (a) and (b) are fundamentally unchanged from (comp), requiring the mappings for composed channels to match. We note $c_{\sigma}(x) = c_{\sigma'}(y) = s$ for all free names x and y since we consider a single session. Clauses (c) and (d) are as in (comp_d), but we remove the condition $d_{\sigma}(z) \neq d_{\sigma'}(y)$ to enable multicut.

THEOREM 8.2. *Let $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma; \mathcal{G}$ and $\Delta = \emptyset$ or Δ contains only 1 or $\perp$. Then: (1) $P \rightarrow^* \mathbf{0}$; and (2) $\text{fuse}(\mathcal{G})$ is a complete global type.*

PROOF. (1) By the same proposition as Proposition 3.6, we obtain if $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$, then $\rho(\sigma(P)) \vdash_{\text{MP}} \Gamma$. By the definition, Γ only contains a single multiparty session where each prefix is simple [31, Definition 5.25 in JACM]. Since $\Delta = \emptyset$ or Δ contains only 1 or $\perp$, Γ is coherent. By the result in Reference [31, Section 5], $\rho(\sigma(P)) \rightarrow^* \mathbf{0}$. Then by the operational correspondence between MP and CLL, $P \rightarrow^* \mathbf{0}$. (2) By construction of (MCut). $\square$

Termination (Theorem 8.2(1)) follows by the fact that the calculus preserves a single (recursion-free) multiparty session [31, Section 5]; The fact that multicut of a complete session generates a complete global type (Theorem 8.2(2)) is by construction. While forms of multicut are in general logically unsound since they invalidate cut elimination [60, Section 6], our (MCut) rule ensures that closed proof terms that may be composed do indeed have progress (and terminate – implying soundness).

We note that in the presence of general delegation (and thus interleaved multiparty sessions) multicut would construct a well-typed (and fuseable) composition of processes that deadlock. However, using the partial type generation in Section 6, we can formulate a simple version of multicut with channel passing, by restricting partial global type generation (Definition 6.7) to only be defined when sent channels are consistently assigned within the same MP session, as defined below.

Definition 8.3 (Multicut with Name Passing). Using the mapping of Section 6 but where all free names are mapped to the same session, by restricting partial global type generation (Definition 6.7) to only be defined when in the clause for channel output, $\forall s \in \sigma. \eta. \langle Q_1 \rangle_{\sigma}^{\eta}(s) = \langle Q_1 \rangle_{\sigma}^{\eta}(\eta(y))$, we define a multicut rule for name passing in Figure 9.

The above (MCut) extends the rule of Figure 8 with clauses (e) and (f) from Figure 7 (red letters highlight differences), which ensure that channels that are to be exchanged are mapped to fresh channels with consistent role assignments in MP.

The same deadlock-freedom as in Theorem 8.2 can be obtained given that the (MCut) rule of Figure 7 stays within a single session.

$$\begin{array}{c}
P \models_{\rho}^{\sigma, \eta} \Delta, x_1:A_1, \dots, x_n:A_n; \mathcal{G}_1 \quad Q \models_{\rho'}^{\sigma', \eta'} \Delta', x_1:A_1^{\perp}, \dots, x_n:A_n^{\perp}; \mathcal{G}_2 \\
\text{fuse}(\mathcal{G}_1, \mathcal{G}_2) \text{ defined } \rho' \cap \rho = \emptyset \quad (\ddagger) \\
\hline
(\text{MCut}) \frac{(\nu x_1, \dots, x_n)(P \mid Q) \models_{\rho''}^{(\sigma' \cup \sigma) \setminus \{x_1, \dots, x_n\}, \eta \cup \eta'} \Delta, \Delta'; \text{fuse}(\mathcal{G}_1, \mathcal{G}_2)}{(\ddagger)}
\end{array}$$

$(\ddagger)$ (a) bound channels: $\rho'' = \rho \cup \rho' \cup_i (x_i, s[p_{\sigma}(x_i)][d_{\sigma}(x_i)])$
(b) role/destination match: $c_{\sigma}(x_i) = c_{\sigma'}(x_i) \quad p_{\sigma}(x_i) = d_{\sigma'}(x_i) \quad d_{\sigma}(x_i) = p_{\sigma'}(x_i)$
(c) unique destination: $\forall y \in \Delta. z \in \Delta'. (d_{\sigma}(y), d_{\sigma'}(z) \notin \rho, \rho' \wedge$
(d) role/destination disjointness: $p_{\sigma}(y) \neq p_{\sigma'}(z) \wedge p_{\sigma}(y) \neq d_{\sigma'}(z) \wedge p_{\sigma'}(z) \neq d_{\sigma}(y)$
(e) bound role/destination match: $\forall x \in \eta. \forall y \in \eta'. x = y \Rightarrow d_{\eta}(x) = p_{\eta'}(y) \vee p_{\eta}(x) = d_{\eta'}(y)$
(f) bound channel disjointness: (1) $\forall x \in \eta. \forall y \in \varphi = \{\sigma', \eta', \rho, \rho'\}. c_{\eta}(x) \neq c_{\varphi}(y)$
(2) $\forall x \in \eta'. \forall y \in \varphi = \{\sigma', \eta, \rho, \rho'\}. c_{\eta'}(x) \neq c_{\varphi}(y)$

Fig. 9. Multicut with channel passing.

We refrain from introducing a multicut for replication since (cut¹) already enables us to share replicated sessions among many clients. Moreover, a cut-free process can only implement a single replicated session (i.e., $P \not\vdash_{\text{CL}} \Xi; x!A, y!B$ cut-free).

Example 8.4 (Two Buyer Protocol [31]). The following defines the coordination of two Buyers seeking to buy from a Seller.

Seller $\triangleq x(t).x\langle 32 \rangle.y\langle 32 \rangle.y.\text{case}\{ok:y(s).0, nok:0\}$
Buyer₁ $\triangleq x\langle \text{"xpto"} \rangle.x(n).z\langle n/2 \rangle.0$
Buyer₂ $\triangleq y(n).z(m).y.ok; y\langle \text{"foo"} \rangle.0.$

The Seller uses channels x and y to interact with Buyer₁ and Buyer₂. Buyer₁ uses channel z to interact with Buyer₂. The Seller waits for Buyer₁ to send it the title of the item that is to be purchased, replying to both buyers with a price. Buyer₁ then emits to Buyer₂ how much it will contribute in the purchase. Buyer₂ then chooses to agree and sends to the Seller an *ok* message and a string (e.g., the shipping address). Consider $\sigma, \sigma_1, \sigma_2$:

$$\begin{aligned}
\sigma(\text{Seller}) &= s[s][b1](t); s[s][b1]\langle 32 \rangle; s[s][b2]\langle 32 \rangle; s[s][b2] \ \& \ \{ok:s[s][b2](s), nok:0\} \\
\sigma_1(\text{Buyer}_1) &= s[b1][s]\langle \text{"xpto"} \rangle; s[b1][s](n); s[b1][b2]\langle n/2 \rangle; 0 \\
\sigma_2(\text{Buyer}_2) &= s[b2][s](n); s[b2][b1](m); s[b2][s] \oplus ok; s[b2][s]\langle \text{"foo"} \rangle; 0.
\end{aligned}$$

In the development of Section 5 the processes above would *not* be composable. We may now compose them via multicut:

$$\begin{aligned}
\langle \text{Seller} \rangle_{\sigma} &= b1 \rightsquigarrow s : \downarrow (\text{str}).s \rightsquigarrow b1 : \uparrow (\text{int}).s \rightsquigarrow b2 : \uparrow (\text{int}). \\
&\quad b2 \rightsquigarrow s : \& \{ok:b2 \rightsquigarrow s : \downarrow (\text{str}).\text{end}, nok:\text{end}\} \\
\langle \text{Buyer}_1 \rangle_{\sigma_1} &= b1 \rightsquigarrow s : \uparrow (\text{str}).s \rightsquigarrow b1 : \downarrow (\text{int}).b1 \rightsquigarrow b2 : \uparrow (\text{int}).\text{end} \\
\langle \text{Buyer}_2 \rangle_{\sigma_2} &= s \rightsquigarrow b2 : \downarrow (\text{int}).b1 \rightsquigarrow b2 : \downarrow (\text{int}).b2 \rightsquigarrow s : \oplus \{ok : b2 \rightsquigarrow s : \uparrow (\text{int}).\text{end}\}.
\end{aligned}$$

Let $G_1 = \langle \text{Seller} \rangle_{\sigma}$, $G_2 = \langle \text{Buyer}_1 \rangle_{\sigma_1}$ and $G_3 = \langle \text{Buyer}_2 \rangle_{\sigma_2}$. Then we have that the fuse of the three global types is as follows:

$$b1 \rightarrow s : (\text{str}).s \rightarrow b1 : (\text{int}).s \rightarrow b2 : (\text{int}).b1 \rightarrow b2 : (\text{int}).b2 \rightarrow s : \{ok : b2 \rightarrow s : (\text{str}).\text{end}\}.$$

Note that we do not require any modifications to the syntax of CLL processes to represent and verify this multiparty example.

Example 8.5 (Multicut and Channel Passing). We illustrate a multicut with channel passing via the following processes:

$$\begin{aligned} P &\triangleq \bar{x}\langle y \rangle. (y\langle 1 \rangle. y(z_1) \mid z\langle 2 \rangle. x(z_2)) \\ Q &\triangleq x(y). y(z_4). x\langle 3 \rangle. y\langle 4 \rangle. w\langle 5 \rangle \\ R &\triangleq z(z_3). w(z_4), \end{aligned}$$

where $P \vdash_{\text{CL}} x : (\text{Int} \otimes \text{Int} \wp 1) \otimes (\text{Int} \wp \perp)$, $z : \text{Int} \otimes 1$, $Q \vdash_{\text{CL}} x : (\text{Int} \wp \text{Int} \otimes \perp) \wp (\text{Int} \otimes 1)$, $w : \text{Int} \otimes \perp$ and $R \vdash_{\text{CL}} z : \text{Int} \wp \perp$, $w : \text{Int} \wp 1$. We define mappings $\sigma, \sigma_1, \sigma_2, \eta, \eta_1, \eta_2$ such that

$$\begin{aligned} \sigma(\eta(P)) &= (\nu s') s[p][q]\langle s'[q_1] \rangle; (s'[p_1][q_1]\langle 1 \rangle; s'[p_1][q_1](z_1) \mid s[p][r]\langle 2 \rangle; s[p][q](z_2)) \\ \sigma_1(\eta_1(Q)) &= s[q][p](y); y[p_1](z_4); s[q][p]\langle 3 \rangle; y[p_1]\langle 4 \rangle; s[q][r]\langle 5 \rangle \\ \sigma_2(\eta_2(R)) &= s[r][p](z_3); s[r][q](z_4). \end{aligned}$$

We then have the following global types:

$$\begin{aligned} \langle P \rangle_{\sigma}^{\eta} &= p \rightsquigarrow q : \uparrow (q_1 \downarrow (\text{Int}); q_1 \uparrow (\text{Int})). p \rightsquigarrow r : \uparrow (\text{Int}). q \rightsquigarrow p : \downarrow (\text{Int}). \text{end} \\ \langle Q \rangle_{\sigma_1}^{\eta_1} &= p \rightsquigarrow q \downarrow (q_1 \downarrow (\text{Int}); q_1 \uparrow (\text{Int})). q \rightsquigarrow p : \uparrow (\text{Int}). q \rightsquigarrow r : \uparrow (\text{Int}). \text{end} \\ \langle R \rangle_{\sigma_2}^{\eta_2} &= p \rightsquigarrow r : \downarrow (\text{Int}). q \rightsquigarrow r : \downarrow (\text{Int}). \text{end}. \end{aligned}$$

We have that $\text{fuse}(\langle P \rangle_{\sigma}^{\eta}, \text{fuse}(\langle Q \rangle_{\sigma_1}^{\eta_1}, \langle R \rangle_{\sigma_2}^{\eta_2})) = p \rightarrow q : (q_1 \downarrow (\text{Int}); q_1 \uparrow (\text{Int})). p \rightarrow r : (\text{Int}). q \rightarrow p : (\text{Int}). q \rightarrow r : (\text{Int}). \text{end}$, and so the multicut rule can be applied, validating the deadlock-free composition of P , Q , and R .

9 DISCUSSION AND RELATED WORK

Interconnectability. Abramsky [1] studies acyclicity of proofs in a computational interpretation of linear logic where names are used exactly once. With session types, channel names can be reused multiple times in a cyclic way (even in CLL defined in Section 2), insofar as two processes may both send and receive along the same channel. This feature, combined with dynamic name creation in CLL, makes the study of interconnectability (and deadlock-freedom in general) more challenging. Abramsky et al. [2] shows that compact-closed categories can interpret cyclic networks typed with multicut rules, but are unable to ensure deadlock-freedom. A categorical model of deadlock-free multicut (i.e., finding additional structure to interpret PMC) is interesting future work. Atkey et al. [3] defines a multicut rule for a variant of CP [60] where propositions are self-dual. This work focuses on capturing the full power of the π -calculus and thus is not concerned with either deadlock-freedom (which their multicut rule does not ensure) or interconnectability.

Multiparty Sessions and Linear Logic. The works of Carbone et al. [11, 14] and Caires and Pérez [7] are the most related to our own. Carbone et al. [14] proposed a typing system for CLL extended to multiparty session primitives. The methodology follows a coherence-based multiparty session type framework, starting from a special form of global types that are annotated by linear logic modalities. The global types are projected to propositions annotated by roles, which type each CLL process. The multicut rule is applied to a complete set of processes in a multiparty session, directly using information of that global type. Carbone et al. [11] develops a variation on the coherence-based approach, based on the works [60] and [7], giving an interpretation of $\otimes$ and $\wp$ that is dual to that of Carbone et al. [14]. We highlight that Carbone et al. [13] also applies a coherence-based linear logic interpretation to a choreographic (global) *language* that differs from multiparty session *types*.

Our work differs from the above approaches in several respects: (1) the works [11, 13, 14] do not (aim to) study interconnectability; (2) the approach of Carbone et al. [11, 14] relies on a special form of global types annotated by modalities, or propositions annotated by participants. Our

compositional PMC-based approach requires no change to the semantics or syntax of processes, types, nor typing rules (except multicut) of CLL; (3) the multicut rule of Carbone et al. [11, 14] is applied to a complete set of processes in a multiparty session, directly using information of global types, while ours is a cut between two processes; and (4) the coherence-based cut rule [11, 14] is more limited than our synthesis-based approach. This is because a complete global type built by PMC characterises all deadlock-free traces observable from local type configurations (Theorem 4.9). For example, our system types Example 6.9, which is untypable by Carbone et al. [11, 14] with the corresponding global type. Thus, our framework allows for more typable representatives of individual multiparty sessions.

Caires and Pérez [7] show that the session interpretation of intuitionistic linear logic can encode the behaviour of multiparty sessions (up to typed barbed congruence). While the goals of our work are quite different, focusing on interconnection networks of CLL processes, we note that their work does not contradict with our results. Their encoding consists of adding *another* participant (i.e., an orchestrator) that mediates all interactions between roles (this encoding also appears in Carbone et al. [11]). Thus, a network such as that of Figure 1(a) is realised by disconnecting all participants, adding a new (fourth) participant p' , and connecting each participant only with p' (i.e., a tree topology). Our encoding preserves the interconnectability of global types, whereas the encoding in the works [7, 11] does not.

Degree of Sharing and Distributability. Dardha and Pérez [19] identify classes of deadlock free processes defined by the number of shared binary sessions between the parallel processes. Our connectability based on multiparty differs from their characterisation and is unrelated to the number of sessions in the process syntax: in our framework, two parallel MP binary processes and CLL have the same connectability since both calculi allow bidirectional interactions ($p \leftrightarrow q$), while in the work [19], CLL is strictly less expressive than two binary session processes with two shared channels. Notice that the work of Dardha and Pérez [19] does not study replication or extensions of CLL with multicut as we have done in Section 7 and Section 8. An encoding criterion for synchronisation among parallel processes called *distributability* is studied in Peters et al. [52]. Their untyped criterion is not applicable to our setting since processes with the same interconnection network might not have the same distributability (and vice versa). For instance, consider

$$\begin{aligned} R &\triangleq s[r][p](y); s'[r_0][q](w); w[q_0](n); y[p_0]\langle n \rangle \\ R' &\triangleq s[r][p](y); y[p_1]\langle n \rangle \mid s'[r_0][q](w); w[q_1](n). \end{aligned}$$

In R , $r \leftrightarrow p$ at s and $r_0 \leftrightarrow q$ at s' are independent, so that R' has the same interconnection structure as R . However, R has 1-distributability while R' has 2-distributability. So, results based on distributability do not in general imply ours.

Progress and Multiparty Compatibility. Type systems for progress in concurrent processes are a vast area of research. See, e.g., the works [18, 34, 50, 59]. While the main emphasis of this work is not a typing system for progress, our encodings ensure progress in restricted interleaved multiparty sessions.

Multiparty Compatibility (MC) properties are studied in the works [22, 37] where a global type is synthesised from communicating automata, assuming *all* participants are initially present. These global synthesis methods are not directly applicable to CLL where composition is binary. To overcome this issue, we proposed PMC together with fusing (partial) global types generated from CLL processes. Investigations of partial compatibility for choreographies [37] and timers [6] would allow us to capture larger classes of connectability (with timing information) in the CLL framework.

A APPENDIX—PROOFS

A.1 Proofs from Section 3—Relating the CLL and MP systems

A.1.1 Proof of Proposition 3.7.

PROPOSITION 3.7 (OPERATIONAL CORRESPONDENCE). *Suppose $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$ and $P \rightarrow P'$. Then $\rho(\sigma(P)) \rightarrow Q$ s.t. $P' \Vdash_{\rho'}^{\sigma'} \Delta'; \Gamma'$ and $Q = \rho'(\sigma'(P'))$ with $\sigma' \subseteq \sigma$, $\rho' \subseteq \rho$.*

PROOF. By induction on the given derivation. Since $\rho(\sigma(P)) \rightarrow Q$, $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$ is derived by applying (comp). Assume the last applied rule is

$$\frac{(\text{comp}) \quad P_1 \vdash_{\text{CL}}^{\sigma} \Delta, x:A \quad P_2 \Vdash_{\rho}^{\sigma'} \Delta', x:A^{\perp}; \Gamma \quad (\star)}{(\nu x)(P_1 \mid P_2) \Vdash_{\rho'}^{(\sigma' \cup \sigma) \setminus \{x\}} \Delta, \Delta'; \Gamma, c_{\sigma}[p_{\sigma}]:\llbracket P_1 \rrbracket_{\sigma}} \quad (\star)$$

and $(\nu x)(P_1 \mid P_2) \rightarrow R$. Consider the case where $\rightarrow$ occurs due to a synchronisation on channel x , hence $R \equiv (\nu x)(P'_1 \mid P'_2)$ with $P'_1 \vdash_{\text{CL}}^{\sigma_1} \Delta, x:A'$ and $P'_2 \Vdash_{\rho_2}^{\sigma_2} \Delta', x:A'^{\perp}; \Gamma'$, where $\sigma_1 \subseteq \sigma$, $\sigma_2 \subseteq \sigma'$, $\Gamma' \subseteq \Gamma$, and $\rho_2 \subseteq \rho \subseteq \rho'$. Hence by weakening the renamings appropriately and applying structural congruence, we have

$$P'_1 \mid P'_2 \Vdash_{\rho_2}^{(\sigma_1 \cup \sigma_2)} \Delta, \Delta'; \Gamma', c_{\sigma_1}[p_{\sigma_1}]:\llbracket P'_1 \rrbracket_{\sigma_1}$$

with $\rho(\sigma((\nu x)(P_1 \mid P_2))) \rightarrow \rho(\sigma((\nu x)(P'_1 \mid P'_2)))$ as required. The cases where either P_1 or P_2 reduce follow by i.h. $\square$

PROPOSITION A.1 (RENAMING). *Suppose $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$. Then:*

- (1) *for all bijective renamings φ on roles and channels, we have $P \Vdash_{\rho \circ \varphi}^{\sigma \circ \varphi} \Delta; \varphi(\Gamma)$.*
- (2) *Assume $P \Vdash_{\rho'}^{\sigma'} \Delta; \Gamma'$. Then there exists a bijective renaming φ on roles and channels such that $\sigma' = \sigma \circ \varphi$ and $\rho' = \rho \circ \varphi$.*

PROOF. By the definition of the mapping. $\square$

Proof of Proposition 3.8.

PROPOSITION 3.8 (THREAD PRESERVATION). *If $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$, then $\rho(\sigma(P))$ is thread preserving.*

PROOF. By induction on $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$. The case (thread) is by Definition 3.1. For the case (comp), by assumption $Q \Vdash_{\rho}^{\sigma'} \Delta', x:A^{\perp}; \Gamma$, $\rho(\sigma'(Q))$ is thread preserving. By the definition of (comp), we note that P is cut-free. Then by (a) in $(\star)$ in (comp), P 's principal participant is p_{σ} . By Definition 3.1 and (b,c) in $(\star)$ in (comp), p_{σ} is disjoint with any principal participant in $\rho(\sigma'(Q))$. Hence the resulting process $\rho'((\sigma \cup \sigma' \setminus \{x\})(\nu x)(P \mid Q))$ is thread preserving. $\square$

Proof of Theorem 3.9.

THEOREM 3.9 (UNIQUENESS). *Assume $P \vdash_{\text{CL}} \Delta$. Suppose $\varphi(P)$ is thread preserving and $\varphi(P)$ is typable by a single MP session, i.e., if $\varphi(P) \vdash_{\text{MP}} \Gamma$ then (1) $\text{dom}(\Gamma)$ contains a single session channel; or (2) $\Gamma = \emptyset$ and $P \equiv \mathbf{0}$. Then there exist ρ and σ such that $\varphi = \sigma \circ \rho$ and $P \Vdash_{\rho}^{\sigma} \Delta; \Gamma$.*

PROOF. By induction on $P \vdash_{\text{CL}} \Delta$. The case $P = \mathbf{0}$ and the case where P is a cut-free process are matched with (thread). By the assumption that a mapping is thread preserving and Proposition A.1, we have that a cut-free process is mapped by (thread) with fixed σ (since (thread) is the only possible rule to create cuts). We show that if σ and ρ are fixed and a process is mapped into a single multiparty session, we must preserve the conditions in $(\star)$ of rule (comp).

The condition (b) of $(\star)$ must hold for P and Q can communicate. The condition $c_{\sigma} = c_{\sigma'}$ ensures a process is mapped into a single session. The condition $d_{\sigma}(z), d_{\sigma'}(y) \notin \rho$ in (c) is required to avoid

a crash of the bound and free names. Finally $d_\sigma(z) \neq d_{\sigma'}(y)$ in (c) is required to map names in Δ and Δ' are mapped to different destinations (participants) to avoid a crash between the originally distinct names in CLL. Notice that without this condition, a parallel composition is untypable in MP since the same role indexed name $s[p][q]$ is spread into two parallel processes. $\square$

A.2 Proofs from Section 4—Partial Multiparty Compatibility

A.2.1 Proof of Theorem 4.9.

THEOREM 4.9 (DEADLOCK-FREEDOM, MC AND EXISTENCE OF A GLOBAL TYPE). *The following are equivalent: (MC) a configuration C is SMC; (DF) C is deadlock-free; (WF) there exists well-formed G such that $\text{Tr}(G) = \text{Tr}(C)$.*

PROOF. (DF) $\Rightarrow$ (MC): (a) Suppose configuration C_0 is deadlock-free. Then by DF, for all $C_0 \xrightarrow{\vec{\ell}}$ C , there exists $C \xrightarrow{\ell_1 \cdot \vec{\ell}_1} C_1 \cdots C_n \xrightarrow{\ell_n \cdot \vec{\ell}_n} C'$ such that C' contains only end.

The base case $C_n = C_0 = C$ is obvious. Suppose $T_{1r} \xrightarrow{\ell'_k} T_{k+1r}$ and we are at C_k . In the case ℓ'_k is an output $\ell'_k = \text{rq} \uparrow (S)$ or a selection $\ell'_k = \text{rq} \triangleleft l$, there are traces such that $C_k \xrightarrow{\vec{\ell}'} C'_k$ that does not include the action from/to the participant p . Hence at C'_k , we have $T_{1r} \xrightarrow{\ell'_k} T_{k+1r}$ and $C'_k \xrightarrow{\ell'_k \cdot \vec{\ell}'_k} C''_k$. This matches with Definition 4.7(1). The case of the input $\ell'_k = \text{rq} \downarrow (S)$ is similar, and matches with Definition 4.7(2). In the case ℓ'_k is a branching such that $\ell'_k = \text{rq} \triangleright l$, we can reach C' which only contains end if and only if there exists $\ell''_k = \text{rq} \triangleright l'$ such that $T_{1r} \xrightarrow{\ell''_k} T'_{k+1r}$ and $C_k \xrightarrow{\vec{\ell}'} C'_k \xrightarrow{\vec{\ell}''_k \cdot \ell''_k} C''_k$. This matches with Definition 4.7(3).

(MC) $\Rightarrow$ (WF) By the synthesis theorem of Deniérou and Yoshida [22].

(WF) $\Rightarrow$ (DF) The trace of well-formed global types is DF by the definition of LTS of the global type G . Then by Proposition 4.6, C is DF. $\square$

A.2.2 Proof of Lemma 4.14.

LEMMA 4.14. *Suppose $\text{fuse}(\text{fuse}(G_i, G_j), G_k)$ with $\{i, j, k\} = \{1, 2, 3\}$ is well formed. Then we have $\text{fuse}(\text{fuse}(G_i, G_j), G_k) \sim_{\text{sw}} \text{fuse}(G_i, \text{fuse}(G_j, G_k))$.*

PROOF. Suppose $\text{fuse}(\text{fuse}(G_1, G_2), G_3)$ is defined. We proceed by the induction of the last rule applied to $\text{fuse}(G_1, G_2)$. The only interesting cases are the first and second rules in Definition 4.10.

Case (1) Let $G_1 = p \leadsto q: \uparrow(T_1).G'_1$ and $G_2 = p \leadsto q: \downarrow(T_2).G'_2$ with $T_1 \geq T_2$. Then by the first fuse rule, we have

$$\text{fuse}(G_1, G_2) = p \rightarrow q: (T_2). \text{fuse}(G'_1, G'_2) = G'_3.$$

Since $\text{fuse}(\text{fuse}(G_1, G_2), G_3)$ is defined, $p \leftrightarrow q \notin G_3$. By applying the third rule,

$$\text{fuse}(G'_3, G_3) = p \rightarrow q: (T_2). \text{fuse}(\text{fuse}(G'_1, G'_2), G_3).$$

We now calculate $\text{fuse}(G_2, G_3)$. Since $p \leftrightarrow q \notin G_3$, by applying the third rule, we have

$$\text{fuse}(G_2, G_3) = p \leadsto q: \downarrow(T_2). \text{fuse}(G'_2, G_3) = G''_3.$$

Then by applying the first rule, we have

$$\text{fuse}(G_1, G''_3) = p \rightarrow q: (T_2). \text{fuse}(G'_1, \text{fuse}(G'_2, G_3)).$$

By induction, $\text{fuse}(\text{fuse}(G'_1, G'_2), G_3) \sim \text{fuse}(G'_1, \text{fuse}(G'_2, G_3))$. Hence, we have $\text{fuse}(\text{fuse}(G_1, G_2), G_3) \sim \text{fuse}(G_1, \text{fuse}(G_2, G_3))$.

Case (2) The case $G_1 = p \rightsquigarrow q : \oplus \{l : G'_1\}$ and $G_2 = p \rightsquigarrow q : \& \{l : G'_2, \{l_j : G_j\}_{j \in J}\}$ is similar to Case (1). Then by the second fuse rule, we have

$$\text{fuse}(G_1, G_2) = p \rightarrow q : \{l : \text{fuse}(G'_1, G'_2)\} = G'_3.$$

Since $\text{fuse}(\text{fuse}(G_1, G_2), G_3)$ is defined, $p \leftrightarrow q \notin G_3$. By applying the fourth rule,

$$\text{fuse}(G'_3, G_3) = p \rightarrow q : \{l : \text{fuse}(\text{fuse}(G'_1, G'_2), G_3)\} = G''_3.$$

We now calculate $\text{fuse}(G_2, G_3)$. Since $p \leftrightarrow q \notin G_3$, by applying the fourth rule, we have

$$\text{fuse}(G_2, G_3) = p \rightsquigarrow q : \& \{l : \text{fuse}(G'_2, G_3), \{l_j : \text{fuse}(G_j, G_3)\}_{j \in J}\} = G'_3,$$

where the well-formedness guarantees $\text{fuse}(G_2, G_3)$ is defined. Then by applying the second rule, we have

$$\text{fuse}(G_1, G'_3) = p \rightarrow q : \{l : \text{fuse}(G_1, \text{fuse}(G'_2, G_3))\}.$$

By induction, $\text{fuse}(\text{fuse}(G'_1, G'_2), G_3) \sim \text{fuse}(G'_1, \text{fuse}(G'_2, G_3))$. Hence, we have $\text{fuse}(\text{fuse}(G_1, G_2), G_3) \sim \text{fuse}(G_1, \text{fuse}(G_2, G_3))$. $\square$

A.2.3 Proof of Theorem 4.15.

THEOREM 4.15 (COMPOSITIONALITY). *Suppose $G_1, \dots, G_n$ are partial global types. Assume $\forall i, j$ such that $1 \leq i \neq j \leq n$, G_i and G_j are PMC and $G = \text{fuse}(G_1, \text{fuse}(G_2, \dots, \text{fuse}(G_n)))$ is a complete global type. Then G is well formed.*

PROOF. By Lemma 4.14, we need only show the case where $G'_{n-1} = \text{fuse}(\dots(\text{fuse}(G_1, G_2), G_3), \dots, G_{n-1})$ and $\text{fuse}(G'_{n-1}, G_n)$ is complete. We proceed by induction on n .

Case $n = 2$. Suppose there are two participants p or q and G_1 contains $p_1 \rightsquigarrow p'_1 \dots p_m \rightsquigarrow p'_m$. Then $p_i = p \wedge p'_i = q$ or $p_i = q \wedge p'_i = p$. Note that we cannot apply the swapping relation between $p_i \rightsquigarrow p'_i$ and $p_j \rightsquigarrow p'_j$ since all actions have the same principal name ($\text{pr}(G_1) = \{p\}$ and $\text{pr}(G_2) = \{q\}$). Then G_2 should contain $p_1 \rightsquigarrow p'_1 \dots p_m \rightsquigarrow p'_m$ with dual modes since the last five rules are not applicable by the side conditions $p \leftrightarrow q \notin G_2$.

Suppose G'_{n-1} contains n participants and there are only partial arrows from p_n or to p_n and $\text{fuse}(G'_{n-1}, G_n)$ is complete. Then the partial arrows in G'_{n-1} form (possibly more than one) chains such that $p_1 \rightsquigarrow p'_1 \dots p_m \rightsquigarrow p'_m$ where either p_i or p'_i is p_n . To obtain a complete global type, we must have dual chains $p_1 \rightsquigarrow p'_1 \dots p_m \rightsquigarrow p'_m$ in G_n . Assuming the completed $n - 1$ participants in G'_{n-1} form a well-formed global type, applying fuse rules one by one from the head (note that the last five rules are not applicable for the partial arrow $p_i \rightsquigarrow p'_i$ in G_1 by the side condition $p_i \leftrightarrow p'_i \notin G_n$), we see that $\text{fuse}(G'_{n-1}, G_n)$ is well formed. $\square$

A.3 Proofs from Section 5—CLL Encoded as a Single Multiparty Session

A.3.1 Proof of Lemma 5.8.

LEMMA 5.8. *Let $P \Vdash_\rho^\sigma \Delta; \Gamma; \mathcal{G}$. $\text{co}(\Gamma)$ implies $\Delta = \emptyset$ or Δ contains only 1 or $\perp$ and $\sigma = \emptyset$.*

PROOF. Assume to the contrary that $\Delta \neq \emptyset$ and does not contain only 1 and $\perp$, or $\sigma \neq \emptyset$. Then it must be the case that P has some free name $x : A \in \Delta$ where $\sigma(x) = s[p][q]$, for some s, p, q with $A \neq 1$ or $\perp$. By construction it must necessarily be the case that $s[p] : T \in \Gamma$. Since x is free in P , we cannot have $s[q] : T' \in \Gamma$ with $T \restriction q \leq \overline{T'} \restriction p$: single thread σ -renamings are invariant on the p role name, so for $s[q]$ to occur in Γ it must have arose due to composition on x , which is impossible since x is free, or on some other (now) bound name y that mapped to $s[q][r]$, for some r . However, since $\sigma(x) = s[p][q]$, by construction we know that $q \neq \rho$. This is contradictory with the assumption of coherence and so we conclude the proof. $\square$

A.3.2 Proof of Lemma 5.9.

LEMMA 5.9. *Let $P \Vdash_{\rho}^{\emptyset} \Delta; \Gamma; \mathcal{G}$ with $\Delta = \emptyset$ or Δ containing only 1 or $\perp$. We have that $\text{co}(\Gamma)$.*

PROOF. Assume to the contrary that Γ is not coherent. Then (1) there exists $s[p]:T, s[q]:T'$ in Γ such that $T \upharpoonright q \not\leq T' \upharpoonright p$ or (2) $s[p]:T \in \Gamma$ such that $q \in \text{roles}(T)$ and $s[q]:T' \notin \Gamma$.

For (1) to be the case, either $T \upharpoonright q$ contains an action unmatched by $T' \upharpoonright p$ or vice versa. Assume wlog that $T \upharpoonright q$ contains an unmatched action. Since both $s[p]:T \in \Gamma$ and $s[q]:T' \in \Gamma$ we know that there exists $(x, s[p][q]) \in \rho$ where some subprocess of P uses $x:A$ for some A and some other subprocess of P uses $\bar{x}:A^{\perp}$. The duality of $x:A$ and $\bar{x}:A^{\perp}$ contradicts the existence of an unmatched action in $T \upharpoonright q$. The argument for an unmatched action in $T' \upharpoonright p$ is identical.

For (2) to be the case, since $s[p]:T \in \Gamma$ and $s[q]:T' \notin \Gamma$, we know that there exists $(x, s[p][r]) \in \rho$, with $r \neq q$, and that $(z, s[q][s]) \notin \rho$. Since $q \in \text{roles}(T)$ then it must be the case that P uses a channel y mapped to $s[p][q]$ that is free, which contradicts our assumptions. $\square$

A.4 Proofs from Section 6: Higher-Order Channel Passing

A.4.1 Proof of Theorem 6.10.

THEOREM 6.10 (CONSISTENCY). *thedeluniqueness Assume $P \vdash_{\text{CL}} \Delta$. If $\varphi(P)$ is typable by Γ , i.e., $\varphi(P) \vdash_{\text{MP}} \Gamma$ and ρ, σ and η satisfy the conditions in Definitions 6.1 and 6.2 and $\varphi = \sigma \circ \eta \circ \rho$, then $P \models_{\rho}^{\sigma, \eta} \Delta; \Gamma; \mathcal{G}$.*

PROOF. By induction on $P \vdash_{\text{CL}}^{\sigma, \eta} \Delta$ (note that $\mathcal{G}$ does not affect the proof, hence we omit). Given bijective renaming $\varphi = \sigma \circ \eta \circ \rho$ that satisfies Definition 6.1 and Definition 6.2, we prove the conditions in (comp_d) are ensured by typability of MP under some σ, η and ρ that satisfy Definition 6.1 and Definition 6.2. The case for conditions (a), (b), (c) is proved as Theorem 6.10. Condition (d) is ensured by the conditions of φ ; and condition (e) is satisfied by the assumption that a delegated bound name coincides with a input variable in the receiver side; and (d) is ensured by a disjointness of free names and bound names in typable MP. $\square$

A.4.2 Proof of Lemma 6.13.

LEMMA 6.13. *Let $P \models_{\rho}^{\emptyset, \eta} \Delta; \Gamma; \mathcal{G}$. $\text{co}(\Gamma)$ implies $\Delta = \emptyset$ or Δ containing only 1 or $\perp$.*

PROOF. We note that the renamings ensure that the two endpoints of an interaction cannot be implemented by the same single-thread process and that bound-names involved in delegation denote linear interactions along different session channels.

Assume to the contrary that $\Delta \neq \emptyset$ and does not contain only 1 and $\perp$, or $\sigma \neq \emptyset$. Then it must be the case that P has some free name $x:A \in \Delta$ where $\sigma(x) = s[p][q]$, for some s, p, q with $A \neq 1$ or $\perp$. By construction it must necessarily be the case that $s[p]:T_1 \in \Gamma$. However, since x is free in P , we have that if $s[q]:T_2 \in \Gamma$ then $T_1 \upharpoonright q \not\leq T_2 \upharpoonright p$. For T_2 to have the corresponding actions with role p there must exist a free name y in P such that $\sigma(y) = s[q][p]$. If both x and y are in the same cut-free sub-process, then this contradicts Definition 6.2. If they are in different sub-processes, then this contradicts the premise of the composition rule. The only remaining possibility is for a bound name of P to have been mapped to $s[q][p]$, which also contradicts the premise of the composition rule. This arguments contradicts the assumption of coherence and so we conclude the proof. $\square$

REFERENCES

- [1] Samson Abramsky. 1993. Computational interpretations of linear logic. *Theoret. Comput. Sci.* 111, 1–2 (1993), 3–57.
- [2] Samson Abramsky, Simon J. Gay, and Rajagopal Nagarajan. 1995. Specification structures and propositions-as-types for concurrency. In *Logics for Concurrency*. 5–40.

- [3] Robert Atkey, Sam Lindley, and J. Garrett Morris. 2016. Conflation confers concurrency. In *A List of Successes That Can Change the World—Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, LNCS, Vol. 9600. Springer, 32–55.
- [4] Gianluigi Bellin and Phil Scott. 1994. On the π -calculus and linear logic. *Theoret. Comput. Sci.* 135, 1 (1994), 11–65.
- [5] Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. 2017. Monitoring networks through multiparty session types. *Theoretical Computer Science* 669 (2017), 33–58.
- [6] Laura Bocchi, Julien Lange, and Nobuko Yoshida. 2015. Meeting deadlines together. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'15)*, Vol. 42. Schloss Dagstuhl, 283–296.
- [7] Luís Caires and Jorge A. Pérez. 2016. Multiparty session types within a canonical binary theory, and beyond. In *Proceedings of the International Conference on Formal Techniques for Distributed Systems (FORTE'16)*, LNCS, Vol. 9688. Springer, 74–95.
- [8] Luís Caires and Frank Pfenning. 2010. Session types as intuitionistic linear propositions. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'10)*, LNCS, Vol. 6269. Springer, 222–236.
- [9] Luís Caires, Frank Pfenning, and Bernardo Toninho. 2016. Linear logic propositions as session types. *Math. Struct. Comp. Sci.* 26, 3 (2016), 367–423.
- [10] Sara Capecchi, Iliara Castellani, and Mariangiola Dezani-Ciancaglini. 2016. Information flow safety in multiparty sessions. *Math. Struct. Comput. Sci.* 26, 8 (2016), 1352–1394.
- [11] Marco Carbone, Sam Lindley, Fabrizio Montesi, Carsten Schürmann, and Philip Wadler. 2016. Coherence generalises duality: A logical explanation of multiparty session types. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'16)*, Leibniz International Proceedings in Informatics, Vol. 59. Schloss Dagstuhl, 33:1–33:15.
- [12] Marco Carbone and Fabrizio Montesi. 2013. Deadlock-freedom-by-design: Multiparty asynchronous global programming. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, (POPL'13)*. 263–274.
- [13] Marco Carbone, Fabrizio Montesi, and Carsten Schürmann. 2014. Choreographies, logically. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'14)*, LNCS, Vol. 8704. Springer, 47–62.
- [14] Marco Carbone, Fabrizio Montesi, Carsten Schürmann, and Nobuko Yoshida. 2015. Multiparty session types as coherence proofs. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'15)*, Leibniz International Proceedings in Informatics, Vol. 42. Schloss Dagstuhl, 412–426.
- [15] Iliara Castellani, Mariangiola Dezani-Ciancaglini, and Paola Giannini. 2017. Concurrent reversible sessions. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'17)*, LIPIcs, Vol. 85. Schloss Dagstuhl, Leibniz-Zentrum fuer Informatik, 30:1–30:17.
- [16] Iliara Castellani, Mariangiola Dezani-Ciancaglini, and Jorge A. Pérez. 2016. Self-adaptation and secure information flow in multiparty communications. *Formal Asp. Comput.* 28, 4 (2016), 669–696.
- [17] Tzu-Chun Chen, Mariangiola Dezani-Ciancaglini, Alceste Scalas, and Nobuko Yoshida. 2017. On the preciseness of subtyping in session types. *Logic. Methods Comput. Sci.* 13, 2 (2017).
- [18] Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. 2016. Global progress for dynamically interleaved multiparty sessions. *Math. Struct. Comput. Sci.* 26, 2 (2016), 238–302.
- [19] Ornela Dardha and Jorge A. Pérez. 2015. Comparing deadlock-free session typed processes. In *Electronic Proceedings in Theoretical Computer Science*, Vol. 190. 1–15.
- [20] Romain Demangeon, Kohei Honda, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. 2015. Practical interruptible conversations: Distributed dynamic verification with multiparty session types and python. *Formal Methods Syst. Des.* 46, 3 (2015), 197–225.
- [21] Romain Demangeon and Nobuko Yoshida. 2015. On the expressiveness of multiparty sessions. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS'15)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 45. Schloss Dagstuhl, 560–574.
- [22] Pierre-Malo Deniérou and Nobuko Yoshida. 2013. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In *International Colloquium on Automata, Languages, and Programming (ICALP'13)*, LNCS, Vol. 7966. Springer, 174–186.
- [23] Henry DeYoung, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2012. Cut reduction in linear logic as asynchronous session-typed communication. In *Computer Science Logic*.
- [24] Luca Fossati, Raymond Hu, and Nobuko Yoshida. 2014. Multiparty session nets. In *Proceedings of the International Symposium on Trustworthy Global Computing (TGC'14)*, LNCS, Matteo Maffei and Emilio Tuosto (Eds.), Vol. 8902. Springer, 112–127.
- [25] Simon Fowler. 2016. An erlang implementation of multiparty session actors. In *Proceedings of the International Conference on Engineering (ICE'16)*, Vol. 223. 36–50.
- [26] S. Gay and M. Hole. 2005. Subtyping for session types in the pi calculus. *Acta Inf.* 42, 2–3 (2005), 191–225.

- [27] Simon J. Gay. 2016. Subtyping supports safe session substitution. In *A List of Successes That Can Change the World—Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*. 95–108.
- [28] J.-Y. Girard. 1987. Linear logic. *Theoret. Comput. Sci.* 50, 1 (1987), 1–102.
- [29] Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. Language primitives and type disciplines for structured communication-based programming. In *European Symposium on Programming (ESOP'98)*, LNCS, Vol. 1381. Springer-Verlag, 22–138.
- [30] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2008. Multiparty asynchronous session types. In *Proceedings of the Conference on Principles of Programming Languages (POPL'08)*. 273–284.
- [31] Kohei Honda, Nobuko Yoshida, and Marco Carbone. 2016. Multiparty asynchronous session types. *J. ACM* 63, 1–9 (2016), 1–67.
- [32] Raymond Hu and Nobuko Yoshida. 2016. Hybrid session verification through endpoint API generation. In *International Conference on Fundamental Approaches to Software Engineering (FASE'16) (LNCS)*, Vol. 9633. Springer, 401–418.
- [33] Raymond Hu and Nobuko Yoshida. 2017. Explicit connection actions in multiparty session types. In *Proceedings of the 20th International Conference on Fundamental Approaches to Software Engineering*, LNCS, Vol. 10202. Springer, 116–133.
- [34] Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniérou, Dimitris Mostros, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. 2016. Foundations of session types and behavioural contracts. *ACM Comput. Surv.* 49, 1, Article 3 (Apr. 2016), 36 pages.
- [35] Limin Jia, Hannah Gommerstadt, and Frank Pfenning. 2016. Monitors and blame assignment for higher-order session types. In *Proceedings of the Conference on Principles of Programming Languages (POPL'16)*. 582–594.
- [36] Julien Lange and Emilio Tuosto. 2012. Synthesising choreographies from local session types. In *Proceedings of the 29th International Conference on Concurrency Theory (CONCUR'12)*, LNCS, Maciej Koutny and Irek Ulidowski (Eds.), Vol. 7454. Springer, 225–239.
- [37] Julien Lange, Emilio Tuosto, and Nobuko Yoshida. 2015. From communicating machines to graphical choreographies. In *Proceedings of the Conference on Principles of Programming Languages (POPL'15)*, Sriram K. Rajamani and David Walker (Eds.). ACM Press, 221–232.
- [38] Julien Lange and Nobuko Yoshida. 2016. Characteristic formulae for session types. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'16)*, LNCS, Vol. 9636. Springer, 833–850.
- [39] Sam Lindley and J. Garrett Morris. 2015. A semantics for propositions as sessions. In *Conference name is European Symposium on Programming (ESOP'15) (LNCS)*, Vol. 9032. Springer, 560–584.
- [40] Sam Lindley and J. Garrett Morris. 2016. Talking bananas: Structural recursion for session types. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming (ICFP'16)*. 434–447.
- [41] Hugo A. Lopez, Eduardo R. B. Marques, Francisco Martins, Nicholas Ng, Casar Santos, Vasco Thudichum Vasconcelos, and Nobuko Yoshida. 2015. Protocol-based verification of message-passing parallel programs. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages & Applications (OOPSLA'15)*. ACM, 280–298.
- [42] Mungo 2016. Mungo homepage. Retrieved from <http://www.dcs.gla.ac.uk/research/mungo/>.
- [43] Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. 2017. Timed runtime monitoring for multiparty conversations. *Formal Asp. Comput.* 29, 5 (2017), 877–910.
- [44] Rumyana Neykova, Raymond Hu, Nobuko Yoshida, and Fahd Abdeljallal. 2018. Session type provider: Compile-time API generation for distributed protocols in F#. (unpublished).
- [45] Rumyana Neykova and Nobuko Yoshida. 2017. Let it recover: Multiparty protocol-induced recovery. In *Proceedings of the 26th International Conference on Compiler Construction*. ACM, 98–108.
- [46] Rumyana Neykova and Nobuko Yoshida. 2017. Multiparty session actors. *Logical Methods Comput. Sci.* 13, 1 (2017).
- [47] Nicholas Ng, Jose Coutinho, and Nobuko Yoshida. 2015. Protocols by default: Safe MPI code generation based on session types. In *International Conference on Compiler Construction (CC'15)*, LNCS, Vol. 9031. Springer, 212–232.
- [48] Nicholas Ng and Nobuko Yoshida. 2016. Static deadlock detection for concurrent go by global session graph synthesis. In *Proceedings of the 25th International Conference on Compiler Construction*. ACM, 174–184.
- [49] Nicholas Ng, Nobuko Yoshida, and Kohei Honda. 2012. Multiparty session C: Safe parallel programming with message optimisation. In *Proceedings of the Technology of Object-Oriented Languages and Systems (TOOLS'12)*, LNCS, Vol. 7304. Springer, 202–218.
- [50] Luca Padovani, Vasco Thudichum Vasconcelos, and Hugo Torres Vieira. 2014. Typing liveness in multiparty communicating systems. In *International Conference on Coordination Models and Languages (COORDINATION) 2014 (LNCS)*, Vol. 8459. Springer, 147–162.
- [51] Jorge A. Pérez, Luís Caires, Frank Pfenning, and Bernardo Toninho. 2012. Linear logical relations for session-based concurrency. In *European Symposium on Programming (ESOP'12)*, LNCS, Vol. 7211. Springer, 539–558.

- [52] Kirstin Peters, Uwe Nestmann, and Ursula Goltz. 2013. On distributability in process calculi. In *European Symposium on Programming (ESOP'13) (LNCS)*, Vol. 7792. Springer, 310–329.
- [53] D. Sangiorgi. 1996. Pi-calculus, internal mobility, and agent passing calculi. *Theor. Comput. Science* 167, 1&2 (1996), 235–274.
- [54] Alceste Scalas, Ornela Dardha, Raymond Hu, and Nobuko Yoshida. 2017. A linear decomposition of multiparty sessions for safe distributed programming. In *Proceedings of the 31st European Conference on Object-Oriented Programming (ECOOP'17)*, Leibniz International Proceedings in Informatics, Vol. 74. Schloss Dagstuhl, Leibniz-Zentrum fuer Informatik, 24:1–24:31.
- [55] Scribble. 2008. Scribble Project. Retrieved from www.scribble.org.
- [56] K. C. Sivaramakrishnan, Mohammad Qudeisat, Lukasz Ziarek, Karthik Nagaraj, and Patrick Eugster. 2013. Efficient sessions. *Sci. Comput. Program.* 78, 2 (2013), 147–167.
- [57] Ramsay Taylor, Emilio Tuosto, Neil Walkinshaw, and John Derrick. 2016. Choreography-based analysis of distributed message passing programs. In *Proceedings of the IEEE/IFIP Workshop on Programmable Data Plane (PDP'16)*. IEEE Computer Society, 512–519.
- [58] Bernardo Toninho, Luís Caires, and Frank Pfenning. 2013. Higher-order processes, functions, and sessions: A monadic integration. In *European Symposium on Programming (ESOP'13) (LNCS)*, Vol. 7792. Springer, 350–369.
- [59] Hugo Torres Vieira and Vasco Thudichum Vasconcelos. 2013. Typing progress in communication-centred systems. In *International Conference on Coordination Models and Language (COORDINATION) 2013 (LNCS)*, Vol. 7890. Springer, 236–250.
- [60] Philip Wadler. 2014. Propositions as sessions. *J. Funct. Program.* 24, 2–3 (2014), 384–418.

Received April 2017; revised March 2018; accepted July 2018