Message-Observing Sessions*

RYAN KAVANAGH, Université du Québec à Montréal, Canada
BRIGITTE PIENTKA, McGill University, Canada

We present Most, a process language with message-observing session types. Message-observing session types extend binary session types with type-level computation to specify communication protocols that vary based on messages observed on other channels. Hence, Most allows us to express global invariants about processes, rather than just local invariants, in a bottom-up, compositional way. We give Most a semantic foundation using *traces with binding*, a semantic approach for compositionally reasoning about traces in the presence of name generation. We use this semantics to prove type soundness and compositionality for Most processes. We see this as a significant step towards capturing message-dependencies and providing more precise guarantees about processes.

CCS Concepts: • **Theory of computation** → **Program specifications**; *Type structures*; Denotational semantics; • **Computing methodologies** → *Concurrent programming languages*.

Additional Key Words and Phrases: session types, program specification, trace semantics, dependent types

1 INTRODUCTION

Session types [Honda 1993; Honda, Vasconcelos, et al. 1998; Takeuchi et al. 1994] allow us to specify and statically verify that processes communicate according to prescribed protocols. Hence, they rule out a wide class of communication-related bugs before executing a given protocol.

A *binary* session type specifies the communication protocol as seen from the point of view of one of the two participants. From a Curry-Howard perspective, it corresponds to the standard sequent calculus proof system for dual intuitionistic linear logic (see [Caires and Pfenning 2010; Caires, Pfenning, and Toninho 2016]), thereby building a logical foundation for specifying and reasoning about concurrent communications. Over the past decade, there have been different approaches to extend it and capture ever richer protocols: value-dependent session types [Toninho, Caires, et al. 2011; Toninho and Yoshida 2018] allow protocols to vary based on previously transmitted values; label-dependent session types [Thiemann and Vasconcelos 2019] describe different communication behaviour depending on labels being observed on a given channel; manifest sharing [Balzer and Pfenning 2017] allows (binary) session types to capture shared communication channels.

In this paper, we describe Most, a language for protocols using *message-observing session types*. A message-observing session type specifies how communication evolves taking into account messages that can be observed on other channels in a process's environment. This is in contrast to both value and label-dependent session types, where a channel's type can only depend on messages that the channel previously carried, but not on messages on other channels in the process's environment. Hence, Most allows us to express global invariants about processes rather than just local invariants.

To specify session types that vary based on messages on other channels, we extend binary session types with *type-level processes*. This allows us to capture a wider range of correctness guarantees than is currently possible. To illustrate, we take a closer look at the identity process id that provides an implementation for the session type $A := \oplus\{\text{left} : \mathbf{1}, \text{right} : \mathbf{1}\}$. This session type specifies that

---

Authors' addresses: Ryan Kavanagh, Département d'informatique, Université du Québec à Montréal, 201 Du Président-Kennedy Avenue, Montréal, QC, H2X 3Y7, Canada, kavanagh.ryan@uqam.ca; Brigitte Pientka, School of Computer Science, McGill University, 3480 University Street, Montréal, QC, H3A 0E9, Canada, bpientka@cs.mcgill.ca.

communication is a message containing either label left or right, followed by a message signalling the end of communication. We can implement id between channels a and b of type A as follows:

$$\text{case } a \{ \text{ left} \Rightarrow b.\text{left}; \text{ wait } a; \text{ close } b \mid \text{right} \Rightarrow b.\text{right}; \text{ wait } a; \text{ close } b \}.$$

Operationally, this process waits to receive a label on a, selects the corresponding branch, and sends the label over b. It then waits for the channel a to close before closing b. However, the session type for channel a and channel b does not capture this precisely. In particular, the type of b does not rule out the following erroneous implementation that swaps labels:

$$\text{case } a \{ \text{ left} \Rightarrow \textcolor{red}{b.\text{right}}; \text{ wait } a; \text{ close } b \mid \text{right} \Rightarrow \textcolor{red}{b.\text{left}}; \text{ wait } a; \text{ close } b \}$$

The crux of the issue is that the messages we wish to allow on b depend on the messages that we observe on a. In Most, we specify such a protocol for b using type-level processes as follows:

$$b : \text{CASE } a \{ \text{ left} \Rightarrow \oplus\{ \text{ left} : \mathbf{1} \} \mid \text{right} \Rightarrow \oplus\{ \text{ right} : \mathbf{1} \} \}.$$

Intuitively, it specifies that the type of b is the unary choice $\oplus\{ \text{ left} : \mathbf{1} \}$ if a label left is observed on a, and symmetrically if right is observed on a. A message-observing session type is necessarily open: it names the channel whose messages it observes. This is analogous to processes, which contain the channel names on which they communicate. The particular channel a observed by the type of b is determined by the process specification in which it appears. For example, the following concrete syntax specifies id by specifying the types of its channels, and it specifies that b's type observes a channel $a : A$ used by id:

$$\text{id} :: \{ \} \; [a : A] \; (b : \text{CASE } a \{ \text{ left} \Rightarrow \oplus\{ \text{ left} : \mathbf{1} \} \mid \text{right} \Rightarrow \oplus\{ \text{ right} : \mathbf{1} \} \}) = \text{case } a \{ \cdots \}.$$

We will revisit this concrete syntax in section 2.

Most provides a "bottom-up" approach to specifying processes and their composition. We still specify the communication protocol from the point of view of one of the two participants, but take into account the messages observed on other channels in the environment. An alternative to specify multi-channel communication patterns are *multiparty* session types. Multiparty session types [Honda, Yoshida, et al. 2016] specify interactions between a static number participants using a *global type*, and project these specifications onto individual participants for typechecking. Despite recent work [Deniélou and Yoshida 2011; Stolze et al. 2023], multiparty session types are an inherently closed world or "top-down" approach, where the entire system must be designed before it can be implemented. In contrast, we can compose protocols in Most from the bottom-up, while still being able to capture some of the richer interactions of multiparty session types.

Concretely, we make the following contributions:

(1) We introduce a **message-observing binary session type system**, called Most. A message-observing session type specifies how communication evolves taking into account communications that can be observed on other channels in our environment. This is achieved using **type-level processes** in session types. These type-level processes further restrict the eligible processes and their actions. We motivate this extension through a sequence of examples.

(2) We give processes and their specifications a **semantics using *traces with binding***, a semantic foundation inspired by nominal sequences [Gabbay and Ghica 2012]. A frequent challenge when we define a trace semantics or ordered semantics for concurrency is that we must account for fresh channel name generation and propagate names involved in higher-order communications. Further, compositionality often requires infinite trace sets or infinite collections of semantic objects, which complicates implementation and obscures the computational content. Instead, we use name binding to quantify over all fresh channel

in a manner reminiscent of higher-order abstract syntax, resulting in a more compact presentation. This abstraction allows us to retain compositionality.

(3) We give a **semantically sound typechecking algorithm**. Process denotations characterize all possible communication behaviours, while specifications specify the processes traces they allow subject to constraints on ambient communications. Soundness ensures that well-typed processes only exhibit behaviours (traces) permitted by their specifications.

We see this work as a significant step towards providing more precise guarantees about processes.

2 MOTIVATION

We introduce the main ideas of MOST through a sequence of examples. This will allow us to showcase both the power of message-observing session types and the design decisions.

2.1 A Quick Guide to Session Types and Specifying Processes

Session types $A, B, \ldots$ specify the communication protocols on named channels $a, b, \ldots$. Processes in our system are organized according to a client-server architecture. This architecture is reflected in our concrete syntax for process specifications:

$$\text{proc} :: \{a_1 : A_1, \ldots, a_n : A_n\} \; [b_1 : B_1, \ldots, b_n : B_n] \; (c_0 : C_0)$$

It states that the process proc is a **server** for (or **provides**) a distinguished service C_0 over a channel c_0. Dually, it is a **client** of (or **uses**) zero or more services $B_1, \ldots, B_n$ on channels named $b_1, \ldots, b_n$, respectively. Finally, its types may refer to zero or more ambient communication channels $a_i : A_i$. Ambient channels are instantiated either by composing proc with processes that use or provide them, or by defining proc as a composition of processes that communicate on these channels. Ambient channels will permit compositional process specifications. To keep our subsequent examples simple, most will feature empty internal and ambient contexts. Each type can refer other channel names in the specification, allowing for mutual dependency between types.

For our first examples, we consider sending a sequence of n bits on a channel a. This is accomplished by sending a sequence of n labels drawn from set $\text{Bit} = \{\, 0, 1 \,\}$. Such a sequence is specified by the session type Bits_n defined by induction on n using the following pair of equations:[1]

$$\text{Bits}_0 = \mathbf{1} \qquad\qquad\qquad \text{Bits}_{n+1} = \oplus\{\, l : \text{Bits}_n \,\}_{l \in \text{Bit}}.$$

The empty sequence of bits is captured by the unit type $\mathbf{1}$: it signals the end of communication on the channel. To define a bit sequence of length $n + 1$, we specify the possible labels (elements) that can be transmitted followed by the sequence of length n. To accomplish this, we use the internal choice type $\oplus\{\, l : \text{Bits}_n \,\}_{l \in \text{Bit}}$.[2] Concretely, we can use the definition to generate the type for bit sequences of length 2 as follows:

$$\text{Bits}_2 = \oplus\{\, 0 : \oplus\{\, 0 : \mathbf{1}, 1 : \mathbf{1} \,\}, 1 : \oplus\{\, 0 : \mathbf{1}, 1 : \mathbf{1} \,\} \,\}$$

It specifies a sequence of two bits, i.e., a sequence of two labels, 0 and 1 followed by termination.

We can then specify a family of identity processes, id_n, that use a channel $a : \text{Bits}_n$ and provide an output channel b of the same type by the following pair of equations:

$$
\begin{aligned}
\text{id}_0 \quad &:: \{\,\} \; [a : \text{Bits}_0 \quad] \; (b : \text{Bits}_0 \quad) \;=\; \text{wait } a; \text{ close } b \\
\text{id}_{n+1} \quad &:: \{\,\} \; [a : \text{Bits}_{n+1}] \; (b : \text{Bits}_{n+1}) \;=\; \text{case } a \; \{\, l \Rightarrow b.l; \; \text{id}_n \,\}_{l \in \text{Bit}}
\end{aligned}
$$

---

[1]MOST omits language-level recursion and relies on meta-level induction to define inductive types. This simplifies MOST's presentation and avoids obscuring its key contributions. We nevertheless sketch how to extend MOST with recursive session types and processes in section 8.4.

[2]An internal choice $\oplus\{\, l : A_l \,\}_{l \in L}$ specifies that a label $l \in L$ will be transmitted, and then communication will satisfy type A_l.

In the base case, id_0, we wait to receive a close message on a that signals the end of communication, and we immediately send it on b before terminating. In the recursive step, id_{n+1}, we perform a case analysis on the label received over a. If we receive a label l, then we take the corresponding branch and send the label l over b. At this point, a and b both have type Bits_n, so we continue as id_n.

2.2 Specifying Identity Processes For Real Using Message-Observing Session Types

The session type specification of the identity process id_n in the previous section fails to ensure that we copy the sequence of bits unchanged from a to b. In Most, we can give a more precise type to the output channel by taking into account the message that we already have received on a given channel using message-observing sessions. In particular, we can refine the type of b to the following more precise family of types $\mathsf{IdBitSeq}_n\ a$ that observe messages on channel a. We give the definition of the identity process next to it, to highlight the close correspondence between the process and the message-observing session type:

Process	Message-Observing Session Type
$\mathsf{id}_0\quad =\ \mathsf{wait}\ a;\ \mathsf{close}\ b$	$\mathsf{IdBitSeq}_0\quad a = \mathsf{CASE}\ a\ \{\ \mathsf{close} \Rightarrow \mathbf{1}\ \}$
$\mathsf{id}_{n+1} = \mathsf{case}\ a\ \{\ l \Rightarrow b.l;\ \mathsf{id}_n\ \}_{l \in \mathsf{Bit}}$	$\mathsf{IdBitSeq}_{n+1}\ a = \mathsf{CASE}\ a\ \{\ l \Rightarrow \oplus\{\ l : \mathsf{IdBitSeq}_n\ a\ \}\ \}_{l \in \mathsf{Bit}}$

In the base case, the type $\mathsf{CASE}\ a\ \{\ \mathsf{close} \Rightarrow \mathbf{1}\ \}$ specifies that we can only signal termination on b after we have observed termination on a. In the recursive step, $\mathsf{IdBitSeq}_{n+1}\ a$ performs a case analysis on the label sent on a. If a label l is sent or received on a, then the type reduces to the unary internal choice $\oplus\{\ l : \mathsf{IdBitSeq}_n\ a\ \}$ that only allows the label l to be sent, and continues as $\mathsf{IdBitSeq}_n\ a$. We can then refine the specification of id to use this more precise protocol:

$$\mathsf{id}_0\quad :: \{\}\ [a : \mathsf{Bits}_0\quad]\ (b : \mathsf{IdBitSeq}_0\quad a) = \mathsf{wait}\ a;\ \mathsf{close}\ b$$
$$\mathsf{id}_{n+1}\quad :: \{\}\ [a : \mathsf{Bits}_{n+1}]\ (b : \mathsf{IdBitSeq}_{n+1}\ a) = \mathsf{case}\ a\ \{\ l \Rightarrow b.l;\ \mathsf{id}_n\ \}_{l \in \mathsf{Bit}}$$

Note that the process and the type of the used channel a remain the same as in the previous section: the only change is to the type of the channel b, which now observes messages on a. At the high-level, the type $\mathsf{IdBitSeq}_n\ a$ ensures that id_n can only output on b what it receives on a.

We call the type $\mathsf{IdBitSeq}_n\ a$ **message observing** because it observes communications on a to determine the range of permitted communications. It can be thought of as a type-level process, as it mimics choices and behaviour specified in the process id_n on the type-level. This type reduces as messages are observed on the channel a. These observations and reductions exist only in the typechecking algorithm (types are not present at runtime), which tracks messages sent or received by the specified process and reduces types accordingly. For example, to type check id_{n+1} against its specification, we check each of its branches "$b.l;\ \mathsf{id}_n$" (to which it steps after receiving a label l) against the specification "$\{\}\ [a : \mathsf{Bits}_n]\ (b : \oplus\{\ l : \mathsf{IdBitSeq}_n\ a\ \})$", which is obtained by reducing all types by the label l received on a.

This example illustrates two key points in Most: (1) message-observing session types can express dependencies on messages on other channels, and (2) message-observing types behave as type-level processes and incorporate a notion of concurrent computation on the type-level.

On a practical level, this allows us to express more precise safety guarantees that cannot be expressed by prior work. On a more theoretical level, we see message-observing session types as an essential step towards a type theory of session types.

2.3 Expressing Mutual Observation in Message-Observing Session Types

Session-typed languages often feature a higher-order session type $A \otimes B$, which specifies a server that sends a channel of type A (call it a) and then communicates according to B.

It is useful in applications for B to observe messages on a. To do so, we introduce the syntax $(a : A) \otimes B$, which binds a name a for the transmitted channel of type A in B. This bound name is eventually instantiated by the actual name of the transmitted channel. This syntax then lets us specify, e.g., sending a pair of identical bit streams: $(a : \mathsf{Bits}_n) \otimes (\mathsf{IdBitSeq}_n\ a)$. A server of this type sends a channel a carrying n bits and then communicating according to the protocol $\mathsf{IdBitSeq}_n\ a$, which specifies transmitting a bit stream equal to the one carried by a.

Communication on both channels occurs independently, so it is natural to ask whether A can also observe messages on the channel of type B (call it b). From the dependent-types perspective in the sequential setting, such mutual dependency may seem unusual, but it captures the concurrent behaviour where the processes providing A and B run concurrently. To allow for mutual observation, we use the syntax $(a : A) \otimes (b : B)$ for higher-order sessions, where name a is bound in B and name b is bound in A. It allows the type B to observe messages on the transmitted channel $a : A$, while also allowing the type A to observe messages on the channel $b : B$. In practice, b will always be instantiated with the name of the channel of type $(a : A) \otimes (b : B)$; we treat it as bound in the type only to ensure that the type makes sense independently of the channel that it types. As a result, we can specify mutually observing higher-order protocols.

To illustrate mutual observation, we consider a load-balanced list service. This service provides a pair of lists, and then forces the client to alternate between lists when receiving elements. It is given by $(a : \mathsf{LBList}_n\ b) \otimes (b : \mathsf{LBList}_n\ a)$, where $\mathsf{LBList}_n\ c$ is inductively defined by:

$$\mathsf{LBList}_0\ c = \mathbf{1} \qquad\qquad \mathsf{LBList}_{n+1}\ c = \oplus\{\, x : \mathsf{CASE}\ c\ \{\, y \Rightarrow \mathsf{LBList}_n\ c\, \}_{y \in \mathsf{Bit}}\, \}_{x \in \mathsf{Bit}}.$$

Intuitively, a channel $a : \mathsf{LBList}_{n+1}\ c$ provides a list element $x \in \mathsf{Bit}$, and then requires the client to observe some element y on c before it can observe another list element from a.

To see how the protocol $(a : \mathsf{LBList}_{n+1}\ b) \otimes (b : \mathsf{LBList}_{n+1}\ a)$ enforces alternation, consider one of its clients. After receiving the channel provided by this service, the client has two channels: channels $a : \mathsf{LBList}_{n+1}\ b$ and $b : \mathsf{LBList}_{n+1}\ a$. Assume, without loss of generality, that it first receives a label on a. The internal choice type specifies that communication on a then satisfies $\mathsf{CASE}\ b\ \{\, y \Rightarrow \mathsf{LBList}_n\ b\, \}_{y \in \mathsf{Bit}}$. In particular, the client cannot communicate on a until it reduces the CASE analysis. To do so, the client must receive a label on b. It can do so, for the type of b is an internal choice $\oplus\{\, x : \mathsf{LBList}_n\ a\, \}_{x \in \mathsf{Bit}}$ (it was reduced by the label received on a). After receiving a label on b, the type of a reduces to $\mathsf{LBList}_n\ b$, while the type of b becomes $\mathsf{LBList}_n\ a$.

For convenience, we write $A \otimes (b : B)$ for $(a : A) \otimes (b : B)$ when a does not appear in B. The abbreviations $(a : A) \otimes B$ and $A \otimes B$ are analogous.

2.4 Depending on Higher-Order Sessions

Protocols can depend on higher-order sessions using an elimination form reminiscent of positive product elimination in functional languages. We illustrate this by specifying an AND logic gate, given by the process AND:

$$\mathsf{AND} :: \{\,\}\ [b : \mathsf{Bits}_1 \otimes \mathsf{Bits}_1]\ (c : \mathsf{Bits}_1) =$$
$$a \leftarrow \mathsf{receive}\ b;\ \mathsf{case}\ a\ \{\, 1 \Rightarrow \mathsf{case}\ b\ \{\, 1 \Rightarrow c.1;\ \mathsf{end} \mid \mathsf{0} \Rightarrow c.\mathsf{0};\ \mathsf{end}\, \}$$
$$\mid \mathsf{0} \Rightarrow \mathsf{case}\ b\ \{\, x \Rightarrow c.\mathsf{0};\ \mathsf{end}\, \}_{x \in \mathsf{Bit}}\, \}$$

where $\mathsf{end} :: \{\,\}\ [a : \mathbf{1}, b : \mathbf{1}]\ (g : \mathbf{1}) = (\mathsf{wait}\ a;\ \mathsf{wait}\ b;\ \mathsf{close}\ g)$. From its specification, we see that AND receives a pair of bits on b and sends a bit on c. Its implementation starts by receiving a channel of type Bits_1 over b and binding it to the name a (syntax: "$a \leftarrow \mathsf{receive}\ b;\ \cdots$"). At this point, the process is a client of $a : \mathsf{Bits}_1$ and $b : \mathsf{Bits}_1$ and provides $c : \mathsf{Bits}_1$. The process then observes the bits on a and b, sends their AND over c, and terminates.

Observe that the above specification does not specify that AND correctly implements an AND gate. However, we can use message observation to specify this functionality at the type level. To do so, we refine the above specification to

$$\text{AND} :: \{\,\} \, [b : \text{Bits}_1 \otimes \text{Bits}_1] \, (c : \text{andOf } b) = \cdots$$

where the following type specifies an AND gate:

$$\text{andOf } b = \text{CASE } b \,\{\, \langle a \rangle \Rightarrow \text{CASE } a \,\{\, 1 \Rightarrow \text{CASE } b \,\{\, 1 \Rightarrow \oplus\{\, 1 : \mathbf{1} \,\} \mid 0 \Rightarrow \oplus\{\, 0 : \mathbf{1} \,\} \,\}$$
$$\mid 0 \Rightarrow \oplus\{\, 0 : \mathbf{1} \,\} \,\} \,\}.$$

This type-level process uses the new construct CASE $b \,\{\, \langle a \rangle \Rightarrow C \,\}$, which binds the name of a channel transmitted over b to a in C. This construct reduces to $[\alpha/a]C$ when a channel α is transmitted on b. This reduction closely mimics the behaviour of the channel-receiving process construct, $a \leftarrow \text{receive } b; P$, which becomes $[\alpha/a]P$ after receiving α on b. In the case of "andOf b", it reduces to the internal choice $\oplus\{\, 1 : \mathbf{1} \,\}$ if the transmitted channel and its carrier both carry a 1 bit, and otherwise it reduces to $\oplus\{\, 0 : \mathbf{1} \,\}$.

2.5 A Fair Auction

In our examples so far, there has been a close correspondence between processes and the message-observing types in their specifications. In fact, this correspondence is partially by design as we take into account the process communication in defining a message-observing type. The next example shows that this correspondence between processes and message-observing types is not always trivial, and that processes can be more complex than the types specifying their communications.

Assume two bidders participating in an auction, where each may bid 0 (described by the label 0) or 1 (described by the label 1). We want to ensure that the bidder with the highest bid wins. With traditional session types, we can specify the protocols for a bidder service as follows, where the external choice $\&\{\cdots\}$ means that the bidder *receives* the label:

$$\text{bidder} = \oplus\{\, 0 : \text{result}, 1 : \text{result} \,\}$$
$$\text{result} = \&\{\, \texttt{lost} : \mathbf{1}, \texttt{tie} : \mathbf{1}, \texttt{win} : \mathbf{1} \,\}$$

A bidder is a process that provides a channel b of type bidder; the name b can freely be renamed, provided that it is kept distinct from other names in the process and its specification:

$$\text{Bidder} :: \{\,\} \, [] \, (b : \text{bidder}) = \cdots$$

An Auctioneer is a client of two bidders and signals the end of the auction on its provided channel c. We can implement it as follows, where $\text{end} :: \{\,\} \, [b_1 : \mathbf{1}, b_2 : \mathbf{1}] \, (c : \mathbf{1}) = (\text{wait } b_1; \text{ wait } b_2; \text{ close } c)$:

$$\text{Auctioneer} :: \{\} \, [b_1 : \text{bidder}, b_2 : \text{bidder}] \, (c : \mathbf{1}) =$$
$$\text{case } b_1 \,\{\, 0 \Rightarrow \text{case } b_2 \,\{\, 0 \Rightarrow b_1.\texttt{tie}; \ b_2.\texttt{tie}; \ \text{end}$$
$$\mid 1 \Rightarrow b_1.\texttt{lost}; b_2.\texttt{win}; \ \text{end} \,\}$$
$$\mid 1 \Rightarrow \text{case } b_2 \,\{\, 0 \Rightarrow b_1.\texttt{win}; \ b_2.\texttt{lost}; \text{end}$$
$$\mid 1 \Rightarrow b_1.\texttt{tie}; \ b_2.\texttt{tie}; \ \text{end} \,\} \,\}$$

Observe that the specification does not rule out an unfair auctioneer that privileges bidder b_1 and always give them the win, i.e., an auctioneer where each branch is $b_1.\texttt{win}; b_2.\texttt{lost}; \text{end}$.

To prohibit unfair auctioneers, we use a message-observing protocol that specifies the interactions for one bidder in terms of the actions of its opponent o:

$$\text{fairAuct } o = \oplus\{\, 0 : \text{CASE } o \,\{\, 0 \Rightarrow \&\{\, \texttt{tie} : \mathbf{1} \,\} \mid 1 \Rightarrow \&\{\, \texttt{lost} : \mathbf{1} \,\} \,\},$$
$$1 : \text{CASE } o \,\{\, 0 \Rightarrow \&\{\, \texttt{win} : \mathbf{1} \,\} \mid 1 \Rightarrow \&\{\, \texttt{tie} : \mathbf{1} \,\} \,\} \,\}$$

The following specification then ensures that Auctioneer is fair:

$$\text{Auctioneer} :: \{\} \; [b_1 : \text{fairAuct } b_2, b_2 : \text{fairAuct } b_1] \; (c : \mathbf{1}) = \cdots.$$

Indeed, the fair auction protocol fairAuct b_2 specifying communications on b_1 ensures that if b_1 sends a label 0, then the auctioneer can only send tie or lost to b_1, and then only after a label 0 or 1 has been observed on b_2, respectively. The case when b_1 bids 1 is analogous.

This example illustrates that specifications do not simply lift process implementations to the type level. In particular, processes can specify more complex interactions than their types. This illustration is even more striking when we consider the following implementation of a bidder. It bids 0, and then terminates after having received the result of the auction:

$$\text{BidsZero} :: \{b_2 : \text{fairAuct } b_1\} \; [] \; (b_1 : \text{fairAuct } b_2) =$$
$$b_1.0; \; \text{case } b_1 \,\{\, \text{result} \Rightarrow \text{close } a \,\}_{\text{result} \in \{\, \texttt{lost}, \texttt{tie}, \texttt{win} \,\}}.$$

This example illustrates how we track the ambient channel $b_2 : \text{fairAuct } b_1$ in the specification to ensure that the type of b_1 remains well-scoped. When typechecking BidsZero against its specification, our typechecking algorithm generates constraints on messages that can appear on b_2 that are then checked when BidsZero is composed with a process implementing b_2.

The ambient context helps ensure that we know the types of all channels observed by types in process interfaces. As a result, we can ensure that types are well-formed relative to each other. It also means that we can modularly specify processes: we specify the types of local channels, but these can observe channels used by other processes in our execution environment.

2.6 Process composition

To illustrate how Most typechecks compositions, we consider the composition of two bit-flipping processes. A bit-flipping process is a client of a bit that provides its negation:

$$\text{neg} :: \{\} \; [i : \text{Bits}_1] \; (o : \text{Bits}_1) = \text{case } i \,\{\, 0 \Rightarrow o.1; \; \text{end} \mid 1 \Rightarrow o.0; \; \text{end} \,\}$$

where end $:: \{\} \; [i : \mathbf{1}] \; (o : \mathbf{1}) = (\text{wait } i; \; \text{close } o)$. We can give neg more precise specifications. For example, we can specify that

$$\text{neg} :: \{\} \; [i : \text{Bits}_1] \; (o : \text{negBit } i) = \cdots$$
$$\text{or} \quad \text{neg} :: \{c : \text{Bits}_1\} \; [i : \text{negBit } c] \; (o : \text{IdBitSeq}_1 \, c) = \cdots,$$
$$\text{where} \quad \text{negBit } a = \text{CASE } a \,\{\, 0 \Rightarrow \oplus\{\, 1 : \mathbf{1} \,\} \mid 1 \Rightarrow \oplus\{\, 0 : \mathbf{1} \,\} \,\}.$$

The first specification ensures that neg sends the negation of the bit it receives on i over o. The second is more interesting: it specifies that if neg receives on i the negation of a bit sent on some ambient channel c, then it outputs that original bit on o.

Using these two specifications, we can ensure that flipping a bit twice behaves as the identity. To do so, we compose an implementation of neg with itself:

$$\text{doubleNeg} :: \{c : \text{negBit } i\} \; [i : \text{Bits}_1] \; (o : \text{IdBitSeq}_1 \, i) = (i \leftrightarrow \text{neg} \leftrightarrow c) \parallel_c (c \leftrightarrow \text{neg} \leftrightarrow o).$$

The syntax $(i \leftrightarrow \text{neg} \leftrightarrow c)$ instantiates the implementation of neg with the channel names i and c, while the syntax $P \parallel_c Q$ composes a server P with its client Q on c. The type of the composition

channel c is determined by the ambient context $\{\,c : \mathsf{negBit}\ i\,\}$. The type checking algorithm checks each composed process against the types specified by the outer specification, i.e., that

$$\mathsf{neg} :: \{\qquad\qquad\} \ [i : \mathsf{Bits}_1\quad] \ (c : \mathsf{negBit}\ i\quad) = \cdots$$
$$\mathsf{neg} :: \{i : \mathsf{Bits}_1\} \ [c : \mathsf{negBit}\ i] \ (o : \mathsf{IdBitSeq}_1\ i) = \cdots.$$

We treat the channel i as ambient in the second specification (it is not accessible to its process). Typechecking the first process poses no difficulties; checking second process is interesting because it is a client of the channel c whose type immediately observes the ambient channel i. To check the second process, we do a case analysis on the messages that could appear on i, reduce the types of c and o accordingly in each case, and check that neg is well typed in the reduced specification. For instance, in the case where $\mathbf{0}$ appears on i, we check that:

$$\mathsf{neg} :: \{i : \mathbf{1}\} \ [c : \oplus\{\,1 : \mathbf{1}\,\}] \ (o : \oplus\{\,\mathbf{0} : \mathbf{1}\,\}) = \cdots.$$

As a side effect, we generate a constraint relating the label $\mathbf{0}$ on i to the behaviour of neg on c and o. After checking both processes independently, our algorithm checks that they impose mutually consistent constraints, and that they satisfy each other's constraints where applicable. In this example, the second process produces a constraint that neg receives 1 on c only if $\mathbf{0}$ appears on i; this constraint is clearly satisfied by the first negation process. Assuming that all constraints are satisfied, we conclude that the composition is well-typed.

The types of composition channels can be locally specified using the program syntax $\nu(a : A)\,.\,P$, which binds a private channel $a : A$ in P. This privacy also ensures that a cannot externally be observed. For example, we can hide the fact that $\mathsf{doubleNeg}$ is implemented as a composition and ensure that no other processes can observe its composition channel c by hiding $c : \mathsf{negBit}\ i$:

$$\mathsf{doubleNeg}' :: \{\,\} \ [i : \mathsf{Bits}_1] \ (o : \mathsf{IdBitSeq}_1\ i) = \nu(c : \mathsf{negBit}\ i)\,.\,\mathsf{doubleNeg}.$$

More generally, the syntax $\nu(a_1 : A_1, \ldots, a_n : A_n)\,.\,P$ binds channels $a_1, \ldots, a_n$ in P. To check $\overrightarrow{\nu a : A}\,.\,P$ against a specification, we check P against the same specification extended with the ambient channels $\overrightarrow{a : A}$. For example, checking $\mathsf{doubleNeg}'$ entails checking the specification $\mathsf{doubleNeg}$.

Our treatment of composition departs from many simply session typed systems [Caires and Pfenning 2010; Wadler 2014] that combine parallel composition with hiding, i.e., that define the composition of P and Q along a as $\nu(a : A)\,.\,P \parallel_a Q$. We distinguish these operations to more flexibly specify process compositions and to ensure that process composition is associative and partially commutative. Indeed, suppose we wished to chain together three processes:

$$P_1 :: \{a_3 : A_3(a_2)\} \ [a_1 : \mathbf{1}] \ (a_2 : A_2(a_3))$$
$$P_2 :: \{\,\} \ [a_2 : A_2(a_3)] \ (a_3 : A_3(a_2))$$
$$P_3 :: \{a_2 : A_2(a_3)\} \ [a_3 : A_3(a_2)] \ (a_4 : \mathbf{1}),$$

where we write $A_i(a_j)$ to mean A_i observes a_j, to form a process $P :: \{\,\} \ [a_1 : \mathbf{1}] \ (a_4 : \mathbf{1})$. The specification of P is clearly well-defined. However, if we always combine parallel composition with hiding, then neither of the following composition specifications is well-defined:

$$P_{12} :: \{\,\} \ [a_1 : \mathbf{1}] \ (a_3 : A_2(a_3)) = \nu(a_2 : A_2(a_3))\,.\,\bigl(P_1 \parallel_{a_2} P_2\bigr)$$
$$P_{23} :: \{\,\} \ [a_2 : A_2(a_3)] \ (a_4 : \mathbf{1}) = \nu(a_3 : A_3(a_2))\,.\,\bigl(P_2 \parallel_{a_3} P_3\bigr)$$

This is because the types $A_2(a_3)$ and $A_3(a_2)$ are ill-scoped after hiding the channels a_2 and a_3. In contrast, if we hide a_2 and a_3 *after* the parallel composition, we can successfully define P as:

$$P :: \{\,\} \ [a_1 : \mathbf{1}] \ (a_4 : \mathbf{1}) = \nu\,(a_2 : A_2(a_3), a_3 : A_3(a_2))\,.\,\bigl(P_1 \parallel_{a_2} \bigl(P_2 \parallel_{a_3} P_3\bigr)\bigr).$$

Our examples show how Most modularly specifies processes while guaranteeing rich invariants.

3 A HIGHER-ORDER PROCESS LANGUAGE

Processes P are generated by the following grammar, where a and b range over channel names, L over sets of choice labels, l and k over labels, and A over types:

Process P, Q ::=	close a	End communication on a and terminate
|	wait a; P	Wait for communication to end on a; continue as P
|	$a.k$; P	Send label k on a; continue as P
|	case $a \{ l \Rightarrow P_l \}_{l \in L}$	Continue as P_l after receiving label $l \in L$ on a
|	send a ($b.P$); Q	Send a channel provided by P on a; continue as Q
|	$b \leftarrow$ receive a; P	Receive channel b on a; continue as P
|	$P \parallel_a Q$	Compose P and Q along channel a
|	$\nu(a_1 : A_1, \ldots, a_n : A_n) . P$	Introduce private channels a_i of type A_i in P

The channel name b is bound in P in $b \leftarrow$ receive a; P and in send a ($b.P$); Q. In send a ($b.P$); Q, it represents the sent channel provided by P. Intuitively, this process forks P with b instantiated by a fresh name, and sends this name over a. The names $a_1, \ldots, a_n$ are bound in P in $\nu(a_1 : A_1, \ldots, a_n : A_n) . P$, and the type annotations exist only to simplify type checking. For the semantics to be reasonable, we require that a does not appear free in P in wait a; P, that P and Q have disjoint sets of free names in send a ($b.P$); Q, and that their free names intersect only in a in $P \parallel_a Q$. As is usual for session-typed processes, we consider *open* processes: the free names in a process name its used and provided channels.

3.1 Trace Semantics

A process denotes a set of traces that describe its possible executions. Traces are sequences of elements $(s; m)$, where m is a message on a channel and s is a tag indicating if m was sent or received, or if it appeared on an unhidden internal channel. To model sending and receiving channels (process constructs send a ($b.P$); Q and $b \leftarrow$ receive a; P), we extend traces with a binding structure similar to the coasbraction operator for nominal sequences [Gabbay and Ghica 2012]. To help ensure that our semantics is compositional, we bind b in the tail of a trace after a channel transmission message. This avoids free name clashes when interleaving traces in the denotation of process composition. Advantageously, our approach keeps the denotations of processes finite, a key property in ensuring that typechecking terminates. Explicitly, traces t, messages m, and signs s are given by the grammar:

Observable Signs o ::= ! | ? | σ	Messages m ::= close on a	Traces t ::= ε
Signs s ::= o | ¿	| label l on a	| $(s;\ m) :: \vec{a} . t$
	| chan on a	

Observable **signs** specify actions performed by processes. They are ? (input), ! (output), and σ (internal synchronization). The constraint sign ¿ is not used in this section, but will be used in the semantics of message-observing types. It specifies a constraint on an ambient channel, i.e., that a message must be observed in the ambient environment before the trace can continue. **Messages** m describe possible messages on a given channel a. The message "close on a" means the end of communication, "label l on a" describes sending a label l, and "chan on a" captures sending a channel. We write $\text{cc}(m)$ for the name of the channel carrying m; in the above cases, $\text{cc}(m) = a$. Finally, **traces** follow a list-like structure. The empty trace is ε. The trace $(s;\ m) :: \vec{a} . t$ prefixes a message m with sign s onto trace t, with zero or more names $\vec{a}$ bound in t. We identify traces up to

α-equivalence. The **free channel names** $\mathrm{fc}(t)$ in a trace t are defined by induction on the structure of t: $\mathrm{fc}(\varepsilon) = \emptyset$ and $\mathrm{fc}((s;\ m) :: \vec{a} . t) = \{\,\mathrm{cc}(m)\,\} \cup (\mathrm{fc}(t) \setminus \vec{a})$.

Processes denote sets of traces, where each trace denotes a possible interleaving of its actions. The denotation $[\![P]\!]$ of a process P is given by induction on its syntax. We explain the semantic clauses and the trace operators they use below:

$$[\![\,\text{close } a\,]\!] = \{\,(!;\ \text{close on } a) :: \varepsilon\,\} \tag{1}$$

$$[\![\,\text{wait } a;\ P\,]\!] = (?;\ \text{close on } a) :: [\![P]\!] \tag{2}$$

$$[\![\,a.k;\ P\,]\!] = (!;\ \text{label } k \text{ on } a) :: [\![P]\!] \tag{3}$$

$$[\![\,\text{case } a \,\{\, l \Rightarrow P_l \,\}_{l \in L}\,]\!] = \bigcup\nolimits_{l \in L} (?;\ \text{label } l \text{ on } a) :: [\![P_l]\!] \tag{4}$$

$$[\![\,\text{send } a\ (b.P);\ Q\,]\!] = (!;\ \text{chan on } a) :: b . ([\![P]\!] \parallel [\![Q]\!]) \tag{5}$$

$$[\![\,b \leftarrow \text{receive } a;\ P\,]\!] = (?;\ \text{chan on } a) :: b . [\![P]\!] \tag{6}$$

$$[\![\,P \,\|_a Q\,]\!] = [\![P]\!] \parallel [\![Q]\!] \tag{7}$$

$$[\![\,\nu(a_1 : A_1, \ldots, a_n : A_n) . P\,]\!] = [\![P]\!] \setminus \{\,a_1, \ldots, a_n\,\} \tag{8}$$

$$\text{where} \quad (s;\ m) :: \vec{a} . T = \{\,(s;\ m) :: \vec{a} . t \mid t \in T\,\}$$

$$T_1 \parallel T_2 = \{\,t \mid t_1 \in T_1,\ t_2 \in T_2,\ t \in (t_1 \parallel t_2)\,\}$$

$$T \setminus \{\,a_1, \ldots, a_n\,\} = \{\,t \setminus \{\,a_1, \ldots, a_n\,\} \mid t \in T\,\}$$

Equation (1) specifies that the only action performed by the process (close a) is sending a close message on a. Dually, eq. (2) specifies that every execution of the process (wait a; P) receives a close message on a before continuing as P.

Equation (3) similarly captures sending a label k on a. Equation (4) specifies that after receiving a label l, the process $P = (\text{case } a \,\{\, l \Rightarrow P_l \,\}_{l \in L})$ continues executing as P_l. The meaning of the process P is then the union of all the traces for P_l prefixed by the received label l.

Equation (5) specifies sending a channel b provided by P over a and then continuing as Q. Processes P and Q execute independently, so the traces that follows sending b are interleavings of the executions of P and Q, i.e., elements of $[\![P]\!] \parallel [\![Q]\!]$. The interleaving operator $\parallel$ is defined in section 5.4. We assume without loss of generality that the bound name b is chosen distinct from any free name in Q. Because P and Q are assumed to have disjoint sets of free channel names, so will their traces. Therefore, $\parallel$ will effectively compute arbitrary interleavings of these two traces. Equation (6) specifies receiving a channel over a and binding it to the name b in P.

Equation (7) states that the parallel composition of processes is given by the synchronized interleavings of their traces. The synchronized interleavings of a pair of traces are defined in such a way that each receive action synchronizes with a send action when available. In particular, it captures a synchronous communication semantics. The process syntax $P\|_a Q$ specifies the channel name a along which P and Q are composed to specify the only name shared between P and Q and to guide typechecking. In contrast, the interleaving operator does not mention a. Our operator is sufficiently general to synchronize traces on multiple channels. This is required, e.g., to handle synchronization for channel transmission, where we must synchronize processes traces on two channels: the transmitted channel and the channel carrying it.

Finally, eq. (8) hides the bound channels a_i by deleting all messages appearing on a_i from the traces in $[\![P]\!]$. It uses the deletion operation $t \setminus \{\,a_1, \ldots, a_n\,\}$ which is given in section 5.1. This deletion operator is defined in such a manner that it deletes not only actions on the a_i, but also actions appearing on the channels they carried, and so on and so forth.

To avoid proliferating definitions for trace operators, we postpone their formal definitions until we have discussed the slightly more general case of traces with constraints; readers wishing to skip ahead are invited to see section 5. Meanwhile, we consider a few intuition-building examples:

Example 3.1. We compute the denotation of the double negation process doubleNeg′ of section 2.6. We start by computing the denotation of the bit flipping process neg composing it:

$$\llbracket i \leftrightarrow \text{neg} \leftrightarrow c \rrbracket = \{ (?;\ \text{label } 0 \text{ on } i) :: (!;\ \text{label } 1 \text{ on } c) :: (?;\ \text{close on } i) :: (!;\ \text{close on } c) :: \varepsilon,$$
$$(?;\ \text{label } 1 \text{ on } i) :: (!;\ \text{label } 0 \text{ on } c) :: (?;\ \text{close on } i) :: (!;\ \text{close on } c) :: \varepsilon \},$$

$$\llbracket c \leftrightarrow \text{neg} \leftrightarrow o \rrbracket = \{ \text{analogous} \}.$$

The denotation of doubleNeg uses the synchronized interleaving operator $\parallel$ to interleave traces from each denotation. Intuitively, it matches $(!; m)$ and $(?; m)$ elements and replaces them by synchronization elements $(\sigma; m)$; the remaining elements are freely interleaved. For example, it interleaves the following traces, respectively from $\llbracket i \leftrightarrow \text{neg} \leftrightarrow c \rrbracket$ and $\llbracket c \leftrightarrow \text{neg} \leftrightarrow o \rrbracket$,

$$(?;\ \text{label } 0 \text{ on } i) :: (!;\ \text{label } 1 \text{ on } c) :: (?;\ \text{close on } i) :: (!;\ \text{close on } c) :: \varepsilon,$$
$$(?;\ \text{label } 1 \text{ on } c) :: (!;\ \text{label } 0 \text{ on } o) :: (?;\ \text{close on } c) :: (!;\ \text{close on } o) :: \varepsilon,$$

to produce, among others, the following traces in $\llbracket (i \leftrightarrow \text{neg} \leftrightarrow c) \parallel_c (c \leftrightarrow \text{neg} \leftrightarrow o) \rrbracket$:

$$(?;\ \text{label } 0 \text{ on } i) :: (\sigma;\ \text{label } 1 \text{ on } c) :: (!;\ \text{label } 0 \text{ on } o) :: (?;\ \text{close on } i) :: (\sigma;\ \text{close on } c) :: \cdots,$$
$$(?;\ \text{label } 0 \text{ on } i) :: (\sigma;\ \text{label } 1 \text{ on } c) :: (?;\ \text{close on } i) :: (!;\ \text{label } 0 \text{ on } o) :: (\sigma;\ \text{close on } c) :: \cdots.$$

Contrasting the two given traces illustrates the concurrent nature of message-passing processes. In the first execution, the left neg process sends 0 before the right receives its close message on i; this order is reversed in the second trace. Finally, the denotation of doubleNeg′ hides all messages on c:

$$\llbracket \nu(c : \text{negBit } i) . \text{doubleNeg} \rrbracket = \{ (?;\ \text{label } 0 \text{ on } i) :: (!;\ \text{label } 0 \text{ on } o) :: (?;\ \text{close on } i) :: \cdots,$$
$$(?;\ \text{label } 0 \text{ on } i) :: (?;\ \text{close on } i) :: (!;\ \text{label } 0 \text{ on } o) :: \cdots, \dots \}.$$

Example 3.2. To illustrate how traces with binding help model channel transmission, consider the parallel composition comp :: $\{\ \}\ [\ ]\ (c : \mathbf{1}) = P \parallel_a Q$, where

$$P :: \{\ \}\ [\ ]\ (a : \mathbf{1} \otimes \mathbf{1}) \qquad = \text{send } a\ (c.\text{close } c);\ \text{close } a$$
$$Q :: \{\ \}\ [a : \mathbf{1} \otimes \mathbf{1}]\ (c : \mathbf{1}) = b \leftarrow \text{receive } a;\ \text{wait } a;\ \text{wait } b;\ \text{close } c$$

Process P denotes two traces, capturing the independent executions of processes close a and close c:

$$\{ (!;\ \text{chan on } a) :: c . (!;\ \text{close on } c) :: (!;\ \text{close on } a) :: \varepsilon,$$
$$(!;\ \text{chan on } a) :: c . (!;\ \text{close on } a) :: (!;\ \text{close on } c) :: \varepsilon \}$$

The process Q exhibits no non-determinism and denotes a single trace:

$$\llbracket Q \rrbracket = \{ (?;\ \text{chan on } a) :: b . (?;\ \text{close on } a) :: (?;\ \text{close on } b) :: (!;\ \text{close on } c) :: \varepsilon \}$$

Their parallel composition denotes a synchronized interleaving of these two trace sets. It is here that binding pays off: we can α-vary the bound name c identifying the transmitted channel in the traces of P to a fresh name b to avoid clashing with the free name c in the trace of Q. The interleaved trace set captures synchronization on both a and the transmitted channel:

$$\llbracket \text{comp} \rrbracket = \{ (\sigma;\ \text{chan on } a) :: b . (\sigma;\ \text{close on } a) :: (\sigma;\ \text{close on } b) :: (!;\ \text{close on } c) :: \varepsilon \}$$

It contains a single trace, because only the first of P's executions successfully synchronizes with Q.

There are no non-terminating executions, i.e., every process performs a finite number of actions:

PROPOSITION 3.3 (TERMINATION). *For all processes P and traces $t \in \llbracket P \rrbracket$, t is finite.*

4 MESSAGE-OBSERVING SESSION TYPES AND PROCESS SPECIFICATIONS

Our type soundness result will show that if a process typechecks against a specification, then its range of behaviours (the traces it denotes) is among those allowed by its specification. We specify which behaviours a process specification allows by giving specifications a trace semantics in section 4.1. First, we define message-observing session types and process specifications.

Recall that each process is a client of some services and the server of a distinguished service. From the perspective of a server of type A, session types are generated by the grammar:

$$
\begin{aligned}
\text{Session Types} \quad A, B ::= \quad &\mathbf{1} &&\text{End communication and terminate} \\
&\mid \ \oplus_{l \in L} A_l &&\text{Internal choice} \\
&\mid \ \&_{l \in L} A_l &&\text{External choice} \\
&\mid \ (a : A) \otimes (b : B) &&\text{Channel transmission} \\
&\mid \ (a : A) \multimap (b : B) &&\text{Channel reception} \\
&\mid \ \text{CASE } a \, \{ \, \text{close} \Rightarrow A \, \} &&\text{Termination observation} \\
&\mid \ \text{CASE } a \, \{ \, l \Rightarrow A_l \, \}_{l \in L} &&\text{Label observation} \\
&\mid \ \text{CASE } a \, \{ \, \langle b \rangle \Rightarrow A \, \} &&\text{Higher-order observation}
\end{aligned}
$$

Session Type Context $\Pi, \Delta, \mathrm{I} ::= a_1 : A_1, \ldots, a_n : A_n$

We explain session types from the perspective of a server. A server of $\mathbf{1}$ sends a close message signalling the end of communication on a channel. The internal choice type $\oplus_{l \in L} A_l$ specifies sending a label $l \in L$, and then communicating according to A_l. Dually, the external choice type $\&_{l \in L} A_l$ specifies receiving a label $l \in L$, and then communicating according to A_l. Many session-typed languages provide a session type $A \otimes B$. A server for a channel c of type $A \otimes B$ sends a channel a of type A on c and then communicates on c according to B. Our syntax $(a : A) \otimes (b : B)$ extends the usual syntax to support mutual observation between the types A and B. It binds the name a in B to allow B to observe communications on the transmitted channel, and it binds b in A to allow A to observe subsequent communications on the carrier channel. A server of a channel c of type $(a : A) \otimes (b : B)$ sends a channel a of type $[c/b]A$, and then communicates on c according to B. Symmetrically, a server of $c : (a : A) \multimap (b : B)$ specifies receiving a channel a of type $[c/b]A$ and then communicating on c according to B. This type has an identical binding structure. We call these types **weak head normal** because they do not involve computation at the outermost level.

Observing types restrict communication based on communication on ambient channels. The type CASE $a \, \{ \, \text{close} \Rightarrow A \, \}$ observes an ambient channel a and reduces to A when a close message is sent on a. The type CASE $a \, \{ \, l \Rightarrow A_l \, \}_{l \in L}$ reduces to the type A_l when a label l is sent on a. Finally, CASE $a \, \{ \, \langle b \rangle \Rightarrow A \, \}$ reduces to A if a channel named b is transmitted over a. Because a priori, we do not know the name of the transmitted channel, CASE $a \, \{ \, \langle b \rangle \Rightarrow A \, \}$ binds the name b in A.

The free channels in a type are inductively defined on the structure in the obvious manner.

Our motivating examples used a concrete syntax "name :: {Σ} [Δ] ($c : C$) = P" to specify a process name, implemented by P, that used channels Δ to provide a service $c : C$, where these channels were free to observe ambient channels Σ. These ambient channels were either (unhidden) internal channels, or channels in other processes whose meaning would be given by process composition. By inspecting the syntax of P, we can partition Σ into a context I of internal channels and a context Π containing the remaining ambient channels. We can then translate the concrete specification syntax {Π, I} [Δ] ($c : C$) into the abstract syntax $\Pi \, \# \, (\Delta \mid \mathrm{I} \vdash c : C)$. Our abstract syntax distinguishes Π and I because they play different roles in typechecking: channels in I can be hidden by P, while those in Π cannot appear in P. Most's typing judgment $P \Vdash \Pi \, \# \, (\Delta \mid \mathrm{I} \vdash c : C) \mathbin{/\!/} \mathfrak{T}$

checks P against its specification, generating constraints $\mathfrak{T}$. Our concrete syntax names processes only to simplify examples; process names do not appear in our formal system.

To specify channel transmission, we need to specify multiple independent processes that execute concurrently. Indeed, send a $(b.P)$; Q spawns process P providing a fresh b, sends b over a, and runs P concurrently with Q providing a. Checking send a $(b.P)$; Q against its specification requires decomposing its specification into specifications for P and Q, and checking P and Q. Accordingly, we extend our abstract syntax for specifications to simultaneously specify multiple independent processes. We write $\Pi \# ((\Delta_1 \mid I_1 \vdash c_1 : C_1) \mid \cdots \mid (\Delta_n \mid I_n \vdash c_n : C_n))$ to simultaneously specify n processes, where channels Π and $\Delta_i, I_i, c_i : C_i$ are ambient in $\Delta_j, I_j, c_j : C_j$ for $i \neq j$. We call $(\Delta \mid I \vdash c : C)$ a process **interface**: it specifies the communication interface for a process. We range over interface sequences $((\Delta_1 \mid I_1 \vdash c_1 : C_1) \mid \cdots \mid (\Delta_n \mid I_n \vdash c_n : C_n))$ using the meta-variable $\mathcal{G}$.

A specification $\Pi \# \mathcal{G}$ is well-defined only if for each $a : A$ in Π or $\mathcal{G}$, the free channels in A are contained in the domain of $\Pi \# \mathcal{G}$, i.e., are assigned types by the specification. This precludes, e.g., $\cdot \# (a : A \mid \cdot \vdash c : \text{CASE } b \{ \text{close} \Rightarrow C \})$ because the channel b is not typed by the specification.

4.1   Type reductions and the semantics of specifications

In our motivating examples, we informally reduced types in specifications after observing messages on channel. To make this explicit, we first define reduction on types, and then extend it to entire specifications. Type reduction is given by a syntactic operation A/π that reduces a type A by an observed communication π. Observed communication, close on a, label l on a, and chan b on a, closely resembles messages from the trace semantics. The only difference is in the case of channel transmissions, where a free channel name b identifies the transmitted channel. This free name will be used by the process typing judgment to update types with the name of the received channel. We write $\text{fc}(\pi)$ for the free channels in π, and $\text{cc}(\pi) = a$ for the **carrier channel**.

Type reduction is partial, capture-avoiding, and inductively defined by:

$$1/\pi = 1$$
$$(\odot_{l \in L} A_l) / \pi = \odot_{l \in L} (A_l / \pi) \qquad\qquad (\odot \in \{ \oplus, \& \})$$
$$((a : A) \odot (b : B)) / \pi = (a : A/\pi) \odot (b : B/\pi) \qquad (\odot \in \{ \otimes, \multimap \})$$

$$(\text{CASE } c \{ \text{close} \Rightarrow A \}) / \pi = \begin{cases} A & \text{if } \pi = \text{close on } c \\ \text{CASE } c \{ \text{close} \Rightarrow A/\pi \} & \text{if } \text{cc}(\pi) \neq c \end{cases}$$

$$(\text{CASE } c \{ l \Rightarrow A_l \}_{l \in L}) / \pi = \begin{cases} A_l & \text{if } \pi = \text{label } l \text{ on } c \text{ and } l \in L \\ \text{CASE } c \{ l \Rightarrow A_l/\pi \}_{l \in L} & \text{if } \text{cc}(\pi) \neq c \end{cases}$$

$$(\text{CASE } c \{ \langle b \rangle \Rightarrow A \}) / \pi = \begin{cases} A & \text{if } \pi = \text{chan } b \text{ on } c \\ \text{CASE } c \{ \langle b \rangle \Rightarrow A/\pi \} & \text{if } \text{cc}(\pi) \neq c \end{cases}$$

All other cases are undefined.

When a message is observed on a channel in a process specification, we simultaneously reduce all types in the specification. This eliminates the need to track past observations and it simplifies the presentation. The syntax $(a_1 : A_1, \ldots, a_n : A_n)/\pi = a_1 : (A_1/\pi), \ldots, a_n : (A_n/\pi)$ captures the simultaneous reduction of types in a context; it is defined whenever all A_i/π are defined. It lifts to process specifications in context-by-context.

A specification $\Pi \# \mathcal{G}$ specifies a collection of permitted behaviours, i.e., a trace set $[\![\Pi \# \mathcal{G}]\!]$. To make this discussion concrete, consider a specification $\Pi \# (\Delta \mid I \vdash c : C)$. The traces in $[\![\Pi \# (\Delta \mid I \vdash c : C)]\!]$ are process traces that involve sending or receiving messages on $\Delta, c : C$, with synchronizations on I. However, these traces may also contain **constraint elements** $(\dot{\iota}; m)$. An

element $(¿; \text{close on } a)$ means that the execution described by the remaining tail of the trace is permitted only if a close message appears on the ambient channel $a \in \text{dom}(\Pi)$. Constraint elements $(¿; \text{label } l \text{ on } a)$ and $(¿; \text{chan on } a)$ are analogous. We use constraint elements to semantically specify that a process is well-typed, provided that some other processes perform some action.

A trace set $[\![\Pi \# \mathcal{G}]\!]$ is the least set satisfying the following collection of inequalities. It is well-defined because type reduction reduces the number of type operators in a specification, so the right-hand specification always contains fewer type operators than the left-hand specification. The right-hand side may be undefined; in this case, we treat it as an empty set.

We first specify label transmission. Where $\pi = \text{label } k \text{ on } a$, $m = \text{label } k \text{ on } a$, and $k \in L$:

$$[\![\Pi \# \mathcal{G} \mid (\Delta \mid I \vdash a : \oplus_{l \in L} A_l)]\!] \supseteq (!; m) :: [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid I/\pi \vdash a : A_k)]\!] \tag{9}$$

$$[\![\Pi \# \mathcal{G} \mid (\Delta, a : \oplus_{l \in L} A_l \mid I \vdash c : C)]\!] \supseteq (?; m) :: [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi, a : A_k \mid I/\pi \vdash c : C/\pi)]\!] \tag{10}$$

$$[\![\Pi \# \mathcal{G} \mid (\Delta \mid I, a : \oplus_{l \in L} A_l \vdash c : C)]\!] \supseteq (\sigma; m) :: [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid I/\pi, a : A_k \vdash c : C/\pi)]\!] \tag{11}$$

The first clause specifies that a provider of type $\oplus_{l \in L} A_l$ can send any label l in L. The provider must obey its original specification, reduced by the just-observed message label l on a. Receiving a label is dually specified. The last clause specifies that a process formed of a composition along $a : \oplus_{l \in L} A_l$ can perform a synchronization on that channel. We define clauses with inequalities to avoid imposing an order on communication on different channels. For example, given a specification $\cdot \# (a : \oplus(0 : 1) \mid \cdot \vdash b : \oplus(0 : 1))$, the above clauses ensure that a and b can be used in either order because its denotation includes traces starting with either channel.

A provider of type 1 can terminate only if it has no other channels. This ensures that channels do not accidentally get discarded and it is captured by the first clause below. Receiving and synchronizing on channels of type 1 ends communication on that channel. The last clause captures termination when there are no processes left. Where $\pi = \text{close on } a$:

$$[\![\Pi \# \mathcal{G} \mid (\cdot \mid \cdot \vdash a : 1)]\!] \supseteq (!; \text{close on } a) :: [\![\Pi/\pi \# \mathcal{G}/\pi]\!] \tag{12}$$

$$[\![\Pi \# \mathcal{G} \mid (\Delta, a : 1 \mid I \vdash c : C)]\!] \supseteq (?; \text{close on } a) :: [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid I/\pi \vdash c : C/\pi)]\!] \tag{13}$$

$$[\![\Pi \# \mathcal{G} \mid (\Delta \mid I, a : 1 \vdash c : C)]\!] \supseteq (\sigma; \text{close on } a) :: [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid I/\pi \vdash c : C/\pi)]\!] \tag{14}$$

$$[\![\Pi \# \cdot]\!] \supseteq \{\varepsilon\} \tag{15}$$

Channel transmission changes the shape of specifications to capture spawning processes. Where $\pi = (\text{chan } b \text{ on } a)$ with b fresh:

$$[\![\Pi \# \mathcal{G} \mid (\Delta_1, \Delta_2 \mid I_1, I_2 \vdash a : (b : B) \otimes (a : A))]\!]$$
$$\supseteq (!; \text{chan on } a) :: b . [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta_1/\pi \mid I_1/\pi \vdash b : B) \mid (\Delta_2/\pi, a : A \mid I_2/\pi \vdash c : C/\pi)]\!] \tag{16}$$

$$[\![\Pi \# \mathcal{G} \mid (\Delta, a : (b : B) \otimes (a : A) \mid I \vdash c : C)]\!]$$
$$\supseteq (?; \text{chan on } a) :: b . [\![\Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi, b : B, a : A \mid I/\pi \vdash c : C/\pi)]\!] \tag{17}$$

The first clause specifies the permitted behaviours of a provider send a $(b.P)$; Q. Operationally, this process forks P and Q, respectively providing b and a, and sends b over a. The used and internal channels of send a $(b.P)$; Q are then divided between P and Q. Accordingly, the first clause splits its interface $\Delta_1, \Delta_2 \mid I_1, I_2 \vdash a : (b : B) \otimes (a : A)$ in two. The channels assigned to P and Q depend on the particular implementation, so the clause allows for arbitrary partitions Δ_1, Δ_2 and I_1, I_2. This clause also illustrates why the observed message chan b on a (used in type reductions) specifies a free name b, in contrast to the message element chan on a used in traces. When we reduce the remaining types in the specification to account for the channel transmission, we must specify the free name b that carries the communications of P. However, this name b is not externally meaningful, and we bind it in the trace to avoid name clashes in process compositions. The second

clause captures receiving a channel over b; it updates the context to contain the received channel. The elided clause for synchronization is analogous to the second clause.

Finally, we specify the constraints generated by ambient channels in Π:

$$[\![\Pi, a : \mathbf{1} \# \mathcal{G}]\!] \supseteq (\text{¿}; \text{ close on } a) :: [\![\Pi/\pi \# \mathcal{G}/\pi]\!], \tag{18}$$

$$\text{where } \pi = \text{close on } a;$$

$$[\![\Pi, a : \odot_{l \in L} A_l \# \mathcal{G}]\!] \supseteq (\text{¿}; \text{ label } l \text{ on } a) :: [\![\Pi/\pi, a : A_l \# \mathcal{G}/\pi]\!], \tag{19}$$

$$\text{where } \pi = \text{label } l \text{ on } a, l \in L, \text{ and } \odot \in \{\oplus, \& \};$$

$$[\![\Pi, a : (b : B) \odot (a : A) \# \mathcal{G}]\!] \supseteq (\text{¿}; \text{ chan on } a) :: b \,.\, [\![\Pi/\pi, a : A, b : B \# \mathcal{G}/\pi]\!], \tag{20}$$

$$\text{where } \pi = \text{chan } b \text{ on } a, \text{ and } \odot \in \{\otimes, \multimap \}.$$

The first clause specifies that if we reduce a specification by observing a close message on an ambient channel a, then the traces permitted by the reduced specification are valid provided that some process in the environment closes a. The second clause is analogous. The third clause captures the fact that sending a channel results in a pair of channels in a context: the transmitted channel and its carrier. We assume without loss of generality that b is fresh in the third clause.

The omitted clauses for $\multimap$ and $\&$ are analogous. There are no clauses for observing types: they allow communication only after reducing to a weak-head normal type. In particular, type reductions are silent and they occur as a result of communication on channels with weak-head normal types.

Example 4.1. The concrete specification neg :: $\{\}$ $[i : \text{Bits}_1]$ $(c : \text{negBit } i)$ of section 2.6 corresponds to the specification $\cdot \# (i : \text{Bits}_1 \mid \cdot \vdash c : \text{negBit } i)$. It only allows the executions

$$(?; \text{ label } 0 \text{ on } i) :: (!; \text{ label } 1 \text{ on } c) :: (?; \text{ close on } i) :: (!; \text{ close on } c) :: \varepsilon,$$

$$(?; \text{ label } 0 \text{ on } i) :: (?; \text{ close on } i) :: (!; \text{ label } 1 \text{ on } c) :: (!; \text{ close on } c) :: \varepsilon,$$

plus the two traces obtained by exchanging 0 and 1. The first two traces illustrate the concurrent nature of communication on different channels. Indeed, negBit i only allows communication on c after communication on i, but no other ordering is imposed between messages on i and c.

Example 4.2. The concrete specification neg :: $\{i : \text{Bits}_1\}$ $[c : \text{negBit } i]$ $(o : \text{IdBitSeq}_1 i)$ of section 2.6 corresponds to the specification $i : \text{Bits}_1 \# (c : \text{negBit } i \mid \cdot \vdash o : \text{IdBitSeq}_1 i)$. It allows:

$$(\text{¿}; \text{ label } 0 \text{ on } i) :: (?; \text{ label } 1 \text{ on } c) :: (!; \text{ label } 0 \text{ on } o) ::$$

$$(\text{¿}; \text{ close on } i) :: (?; \text{ close on } c) :: (!; \text{ close on } o) :: \varepsilon,$$

$$(\text{¿}; \text{ label } 0 \text{ on } i) :: (\text{¿}; \text{ close on } i) :: (?; \text{ label } 1 \text{ on } c) :: (!; \text{ label } 0 \text{ on } o) ::$$

$$(?; \text{ close on } i) :: (!; \text{ close on } o) :: \varepsilon, \text{ etc.}$$

Interpreting each trace as an execution or sequence of actions permitted to neg, the first trace specifies that if a label 0 is sent on an ambient channel i, then neg can receive a bit 1 on c, send 0 on o, and so on and so forth. The second trace is analogous, but illustrates that we impose no extraneous orderings on messages on different channels.

Example 4.3. The specification $a : \mathbf{1}, b : \mathbf{1} \# (\cdot \mid \cdot \vdash c : \text{CASE } a \, \{ \text{close} \Rightarrow \text{CASE } b \, \{ \text{close} \Rightarrow \mathbf{1} \} \})$ permits two traces:

$$(\text{¿}; \text{ close on } a) :: (\text{¿}; \text{ close on } b) :: (!; \text{ close on } c) :: \varepsilon,$$

$$(\text{¿}; \text{ close on } b) :: (\text{¿}; \text{ close on } a) :: (!; \text{ close on } c) :: \varepsilon.$$

These two traces illustrate that the type of c does not impose an order on ambient communications: because we can reduce under CASE, communication on a and b can occur in any order. Indeed, the first trace corresponds to reducing the specification to $b : \mathbf{1} \# (\cdot \mid \cdot \vdash c : \text{CASE } b \, \{ \text{close} \Rightarrow \mathbf{1} \})$

(after observing close on a) and then to $\cdot \# (\cdot \mid \cdot \vdash c : \mathbf{1})$ (after observing close on b); the second captures first reducing to $a : \mathbf{1} \# (\cdot \mid \cdot \vdash c : \text{CASE } a \{ \text{close} \Rightarrow \mathbf{1} \})$, then to $\cdot \# (\cdot \mid \cdot \vdash c : \mathbf{1})$.

Example 4.4. The specification $\cdot \# (\cdot \mid \cdot \vdash a : (b : \mathbf{1}) \otimes \text{CASE } b \{ \text{close} \Rightarrow \mathbf{1} \})$ specifies sending a channel of type $\mathbf{1}$, and with the carrier closing after the sent channel. Its trace set is:

$$(!; \text{ chan on } a) :: b \, . \, [\![\cdot \# (\cdot \mid \cdot \vdash b : \mathbf{1}) \mid (\cdot \mid \cdot \vdash a : \text{CASE } b \{ \text{close} \Rightarrow \mathbf{1} \})]\!]$$

$$= (!; \text{ chan on } a) :: b \, . \, (!; \text{ close on } b) :: [\![\cdot \# (\cdot \mid \cdot \vdash a : \mathbf{1})]\!]$$

$$= \{ \, (!; \text{ chan on } a) :: b \, . \, (!; \text{ close on } b) :: (!; \text{ close on } a) :: \varepsilon \, \}.$$

This example shows how multiple interfaces in a specification are used to specify relationships between messages from different processes that are spawned in the course of channel transmission.

5 OPERATIONS ON TRACES

Our trace semantics for processes uses two operations on traces: deletion and synchronized interleaving. Name deletion deletes all messages in a trace whose names appear in a given set, while synchronized interleavings interleave two traces while ensuring that input and output actions match up. We also specify how to delete constraints elements from traces. All three operations will be used by our typechecking algorithm when generating traces with constraints.

5.1 Deleting channel names from traces

The construct $\nu(\overrightarrow{a_i : A_i}) \, . \, P$ binds the channels a_i to hide them from external view. Semantically, hiding is captured by deleting all messages on channels a_i from the traces of P. To account for channel transmission, we may need to delete messages on other channel names. In particular, if a_i carries some channel b, then we must also delete all messages on b from the trace.

Given a set X of names and a trace t, the deletion of X in t is the trace $t \setminus X$ inductively defined by:

$$\varepsilon \setminus X = \varepsilon \quad \text{and} \quad ((s; \, m) :: c_1, \ldots, c_n \, . \, t) \setminus X = \begin{cases} t \setminus (X \cup \{ c_1, \ldots, c_n \}) & \text{if } \mathrm{cc}(m) \in X \\ (s; \, m) :: c_1, \ldots, c_n \, . \, (t \setminus X) & \text{otherwise} \end{cases}$$

We assume without loss of generality that the names c_i are chosen distinct from those already in X.

Example 5.1. If $t = (?; \text{ chan on } a) :: b \, . \, (?; \text{ close on } a) :: (?; \text{ close on } b) :: (!; \text{ close on } c) :: \varepsilon$, then $t \setminus \{ a \} = (!; \text{ close on } c) :: \varepsilon$.

5.2 Deleting constraints from traces

The operation $t/¿$ deletes all constraints from t. We will use it when checking that processes have consistent constraint sets when typing process compositions. It is inductively defined by

$$\varepsilon/¿ = \varepsilon \qquad \text{and} \qquad (s; \, m) :: \vec{x} \, . \, t/¿ = \begin{cases} t/¿ & \text{if } s = ¿ \\ (s; \, m) :: \vec{x} \, . \, t/¿ & \text{otherwise} \end{cases}$$

5.3 Trace reduction

When interleaving two traces, we need to delete constraints in one trace satisfied by messages in the other trace. We do so using a partial operation t/π that deletes a constraint $(¿; m)$ in t if it is satisfied by the observed communication π. It is inductively defined on the structure of the trace:

$$\varepsilon/\pi = \varepsilon \qquad \qquad ((¿; \text{ close on } a) :: t)/(\text{close on } a) = t$$

$$((¿; \text{ label } l \text{ on } a) :: t)/(\text{label } l \text{ on } a) = t \qquad ((¿; \text{ chan on } a) :: c \, . \, t)/(\text{chan } b \text{ on } a) = [b/c]t$$

$$((s; \, m) :: \vec{x} \, . \, t)/\pi = (s; \, m) :: \vec{x} \, . \, (t/\pi) \quad \text{if } \mathrm{cc}(m) \neq \mathrm{cc}(\pi),$$

and it is undefined in all other cases. The fourth clause instantiates a bound name c for a transmitted channel with the free name b actually observed in the course of channel transmission. The last clause handles an observation on a channel different from the head of the trace. It is defined only if t/π is defined. We always α-vary $\vec{x}$ to be distinct from the names in π.

Example 5.2. The reduction $((\mathbf{\imath};\ \text{label } k \text{ on } a) :: \varepsilon)/(\text{label } l \text{ on } a)$ is undefined because the label l observed on a does not match the label k expected by the constraint. Similarly, the reduction $((!;\ \text{close on } a) :: \varepsilon)/(\text{label } l \text{ on } a)$ is undefined because the trace we are reducing sends on a, while the observation "label l on a" is meant to be an observation on an ambient channel a.

5.4 Synchronized trace interleavings

Synchronized interleavings interleave traces according to a synchronous communication semantics. It marks matching input-output elements as synchronized, and ensures that constraints imposed by one trace are not violated by the other trace. This commutative operator is defined by lexicographic induction on the length of the traces. We group its clauses into four categories:

(1) the base cases: interleaving a trace with an empty trace;
(2) synchronizing cases: the heads of each trace perform actions on the same channel;
(3) commuting cases: the heads of each trace perform actions on different channels;
(4) ill-defined cases: unsynchronizeable sends and receives, etc.

The base cases specify that the empty trace is the unit for interleaving:

$$\varepsilon \parallel t = \{\,t\,\}, \qquad t \parallel \varepsilon = \{\,t\,\}.$$

The first synchronizing case holds whenever two traces send and receive the same message:

$$((!;\ m) :: \vec{x} \,.\, t_1) \parallel ((?;\ m) :: \vec{x} \,.\, t_2) = (\sigma;\ m) :: \vec{x} \,.\, (t_1 \parallel t_2) \,.$$

It specifies that these two traces synchronize (captured by the σ sign), and that the tail is given by interleaving their tails. To ensure that the resulting trace is well-formed, we require that both traces bind the same names. It is this implicit α-variation that matches sent and received channel names b when interleaving traces of the form $(!;\ \text{chan on } a) :: b \,.\, t_1$ and $(?;\ \text{chan on } a) :: b \,.\, t_2$.

The other principal case manages constraints. It states that a constraint in one trace can be dropped whenever it is satisfied by a corresponding action in the other trace, or whenever both traces have the same constraint. If a constraint is matched against a different message on that channel, then the constraint cannot be satisfied and there are no interleavings.

$$((s;\ m) :: \vec{x} \,.\, t_1) \parallel ((\mathbf{\imath};\ m) :: \vec{x} \,.\, t_2) = (s;\ m) :: \vec{x} \,.\, (t_1 \parallel t_2)$$

The commuting cases handle heads that act on different channels. Let $(m \propto \vec{x})$ be the observed communication induced by a message m and channel names $\vec{x}$: $(\text{close on } a \propto \emptyset) = \text{close on } a$; $(\text{label } l \text{ on } a \propto \emptyset) = \text{label } l \text{ on } a$; $(\text{chan on } a \propto b) = \text{chan } b \text{ on } a$; and $(m \propto \vec{x})$ is undefined otherwise. Given a partial map f, let $f(x) \simeq y$ define y whenever $f(x)$ is defined. Where $T_i = (s_i;\ m_i) :: \vec{x}_i \,.\, t_i$,

$$\left. \begin{aligned} &(s_1;\ m_1) :: \vec{x}_1 \,.\, t_1 \parallel (s_2;\ m_2) :: \vec{x}_2 \,.\, t_2 \\ &= \{\, (s_1;\ m_1) :: \vec{x}_1 \,.\, t \mid T_2/(m_1 \propto \vec{x}_1) \simeq T_2' \wedge t \in t_1 \parallel T_2' \,\} \cup \\ &\cup \{\, (s_2;\ m_2) :: \vec{x}_2 \,.\, t \mid T_1/(m_2 \propto \vec{x}_2) \simeq T_1' \wedge t \in T_1' \parallel t_2 \,\} \end{aligned} \right\} \quad \text{if } \mathrm{cc}(m_1) \neq \mathrm{cc}(m_2).$$

This case is well-defined because the lengths of T_1' and T_2' are respectively at most the lengths of T_1 and T_2. Our trace reduction operator hides considerable complexity, so we unpack the definition.

First, observe that the commuting case ensures that if a synchronization between $(s_1; m_1)$ and an element of T_2 is eventually possible, then it occurs. Indeed, if $\mathrm{cc}(m_1)$ appears free in T_2 with an observable sign, then $T_2/(m_1 \propto \vec{x}_1)$ is undefined (see, e.g., example 5.2). As a result, the first set

(prefixing $(s_1; m_1)$) is empty, and we interleave T_1 with the tail of T_2 in the second set to eventually synchronize $(s_1; m_1)$ with the corresponding element of T_2.

Example 5.3. The interleaving $(!; \text{ close on } a) :: \varepsilon \parallel (?; \text{ close on } b) :: (?; \text{ close on } a) :: \varepsilon$ is determined by the commuting case. The reduction in the first set of the clause is undefined, so the first set is empty. The second, $\{ (?; \text{ close on } b) :: t \mid t \in ((!; \text{ close on } a) :: \varepsilon) \parallel ((?; \text{ close on } a) :: \varepsilon) \}$, is defined in terms of a principal case. This principal case captures the synchronization on a, and the interleavings are given by $\{ (?; \text{ close on } b) :: (\sigma; \text{ close on } a) :: \varepsilon \}$.

Second, the commuting case detects deadlock. Indeed, if two traces attempt to synchronize on different channels in opposite orders, then the carrier of the head of one trace is free in the other trace. This implies that the reduction in each set is undefined, resulting in no interleavings:

$$((?; \text{ close on } a) :: (!; \text{ close on } b) :: \varepsilon) \parallel ((?; \text{ close on } b) :: (!; \text{ close on } a) :: \varepsilon) = \emptyset.$$

Third, the commuting case uses reduction to handle constraint satisfaction. Indeed, if $(s_1; m_1)$ satisfies a constraint in T_2, then the reduction $T_2 / (m_1 \propto \vec{x}_1)$ deletes that constraint before interleaving the result with the tail t_1.

Example 5.4. Interleaving $((!; \text{ close on } a) :: \varepsilon) \parallel ((?; \text{ close on } b) :: (¿; \text{ close on } a) :: \varepsilon)$ produces the interleavings $\{ (!; \text{ close on } a) :: (?; \text{ close on } b) :: \varepsilon, \quad (?; \text{ close on } b) :: (!; \text{ close on } a) :: \varepsilon \}$, where the send on a in the first trace satisfies the constraint in the second trace.

Finally, the commuting case ensures that constraint satisfaction obeys synchronization boundaries. For example, consider interleaving the traces $(¿; \text{ close on } a) :: (?; \text{ close on } b) :: \varepsilon$ and $(!; \text{ close on } b) :: (!; \text{ close on } a) :: \varepsilon$. Intuitively, they synchronize on b. However, the constraint on a cannot be satisfied by the send on a, because the constraint is imposed before the synchronization, while the send occurs after the synchronization on b. This unsatisfiability implies that the traces should have no synchronizing interleavings. This is captured by our interleaving operator: the only applicable clause is the commuting case, and both sets defining it are empty.

All other cases are taken to be ill-defined, so we define their set of interleavings to be empty. These include, e.g., cases of the form $((!; \text{ close on } a) :: t) \parallel ((!; \text{ label } l \text{ on } a) :: t')$ (sending on the same channel) and mismatched sends and receives: $((!; \text{ close on } a) :: t) \parallel ((?; \text{ label } l \text{ on } a) :: t')$.

A trace is **safely constrained** if any carrier channel associated with a constraint sign is not also associated with input, output, or synchronization. By proposition 5.5, constraints can only limit the interleavings of observable actions. In other words, checking constraints when interleaving traces does not introduce new observable behaviours. The opposite inclusion will play a key role when typechecking process composition:

PROPOSITION 5.5. *Given safely constrained traces t_1 and t_2, $(t_1 \parallel t_2) / ¿ \subseteq (t_1 / ¿) \parallel (t_2 / ¿)$.*

5.5 Type-theoretic properties

Section 4.1 gave a trace semantics for process specifications. We relate the behaviour of trace operators to this trace assignment. Recall that a process specification $\Pi \# \mathcal{G}$ is well-formed only if every channel appearing free in a type in $\Pi, \mathcal{G}$ is typed by some other type assignment in the specification.

Proposition 5.6 semantically characterizes hiding internal channels in a specification. The well-formedness hypothesis is required to ensure that no types become ill-scoped. For example, it precludes deleting $a : \mathbf{1}$ from $\Pi \# (\Delta \mid I, a : \mathbf{1} \vdash c : \text{CASE } a \{ \text{close} \Rightarrow C \})$.

PROPOSITION 5.6. *If $t \in \llbracket \Pi \# \mathcal{G} \mid (\Delta \mid I, \Gamma \vdash a : A) \rrbracket$ and $\Pi \# \mathcal{G} \mid (\Delta \mid I \vdash a : A)$ is well-formed, then $(t \setminus \text{dom}(\Gamma)) \in \llbracket \Pi \# \mathcal{G} \mid (\Delta \mid I \vdash a : A) \rrbracket$.*

The following proposition characterizes independent parallel executions. To concretize it, consider specifications $\Pi, \Delta_2, I_2, a_2 : A_2 \# (\Delta_1 \mid I_1 \vdash a_1 : A_1)$ and $\Pi, \Delta_1, I_1, a_1 : A_1 \# (\Delta_2 \mid I_2 \vdash a_2 : A_2)$. They specify two processes whose types can observe each other's channels. We can compose these two specifications in a single specification: $\Pi \# (\Delta_1 \mid I_1 \vdash a_1 : A_1) \mid (\Delta_2 \mid I_2 \vdash a_2 : A_2)$. The proposition states that interleaving the traces of each individual specification results in traces allowed by the composed specification. This proposition is useful when characterizing traces involving channel transmissions. More generally, where $\overline{\mathcal{G}}$ denotes the typed channels in $\mathcal{G}$:

PROPOSITION 5.7. *If* $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \# \mathcal{G}_1]\!]$ *and* $t_2 \in [\![\Pi, \overline{\mathcal{G}_1} \# \mathcal{G}_2]\!]$, *then* $t_1 \parallel t_2 \subseteq [\![\Pi \# \mathcal{G}_1 \mid \mathcal{G}_2]\!]$.

An alternate form of specification composition results in synchronization. Indeed, consider processes specifications $\Pi, \Delta_2, I_2, c : C \# (\Delta_1 \mid I_1 \vdash a : A)$ and $\Pi, \Delta_1, I_1 \# (\Delta_2, a : A \mid I_2 \vdash c : C)$. The parallel composition operation composes them along the channel a so that they synchronize on this channel. Subject to side conditions, their composition satisfies $\Pi \# (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C)$ where a becomes an inner synchronization channel. This synchronizing parallel composition is also characterized by synchronized interleaving:

PROPOSITION 5.8. *If* $t_1 \in [\![\Pi, \Delta_2, I_2, c : C \# (\Delta_1 \mid I_1 \vdash a : A)]\!]$ *and* $t_2 \in [\![\Pi, \Delta_1, I_1 \# (\Delta_2, a : A \mid I_2 \vdash c : C)]\!]$, *then* $t_1 \parallel t_2 \subseteq [\![\Pi \# (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C)]\!]$.

6 TYPECHECKING PROCESSES

We give a typechecking algorithm to statically check that a process satisfies its specification, i.e., that each execution of the process is permitted by its specification. It is driven by a collection of inference rules and operates in two alternating phases.

The first phase decomposes process compositions into individual processes. It then determines the next channel a that a given process will use to communicate, and kicks off the second phase. The second phase focusses on the type of a. It first tries to reduce the focussed type to weak head normal form, and then proceeds to typecheck the given process with respect to that weak head normal type. Because our system is compositional, types may depend on ambient channels whose meaning will eventually be given by composition, and type reduction may cause constraints to be imposed on ambient communications. To capture these, our algorithm generates a collection constraints that must be satisfied by the environment in order for the process to be well-typed.

6.1 Phase one: uniform process typing

The first phase, which we call uniform typing, is driven by a judgment $P \Vdash \Pi \# (\Delta \mid I \vdash a : A) /\!\!/ \mathfrak{T}$ inductively defined by the rules of fig. 1. We use colours to indicate the modes of each parameter in the judgment, where red indicates inputs to the judgment and blue indicates outputs. Here, we attempt to check process P against the well-formed specification $\Pi \# (\Delta \mid I \vdash a : A)$. The judgment outputs a set $\mathfrak{T}$ of traces that consist of the communications of P interleaved with constraints on channels in Π. By our soundness result (theorem 7.1), restricting each trace in $\mathfrak{T}$ to its observable actions (i.e., those with sign ?, !, or σ) will result in the set $[\![P]\!]$ of traces of P. We implicitly assume that process specifications are well-formed, i.e., that all names are in scope.

The rule (NEW) introduces internal channels that can be used to type process P. The constraint set for its conclusion is obtained by deleting the new channels from the constraint set for P. This is analogous to the trace semantics of hiding in section 3.1. Because we assume that all specifications are well-formed, we know that no types in $\Pi \# (\Delta \mid I \vdash c : C)$ attempt to observe $a_1, \ldots, a_n$.

The rule (PAR) ensures that two processes can be composed only if they impose compatible constraints. This is ensured by premise $\mathfrak{T}_1 \therefore \mathfrak{T}_2$ that holds if and only if:

$$\forall t_1 \in \mathfrak{T}_1 \;.\; \forall t_2 \in \mathfrak{T}_2 \;.\; (t_1/_{\dot{\iota}} \parallel t_2/_{\dot{\iota}}) = (t_1 \parallel t_2)/_{\dot{\iota}}. \tag{21}$$

$$\boxed{P \Vdash \Pi \ \# \ (\Delta \mid I \vdash a : A) \ /\!\!/ \ \mathfrak{T}}$$ Process P satisfies specification $\Pi \ \# \ (\Delta \mid I \vdash a : A)$ subject to constraints $\mathfrak{T}$

$$\frac{P \Vdash \Pi \ \# \ (\Delta \mid I, a_1 : A_1, \ldots, a_n : A_n \vdash c : C) \ /\!\!/ \ \mathfrak{T}}{\nu(a_1 : A_1, \ldots, a_n : A_n) . P \Vdash \Pi \ \# \ (\Delta \mid I \vdash c : C) \ /\!\!/ \ \mathfrak{T} \setminus \{a_1, \ldots, a_n\}} \ (\text{New})$$

$$\frac{P \Vdash \Pi, \Delta_2, I_2, c : C \ \# \ (\Delta_1 \mid I_1 \vdash a : A) \ /\!\!/ \ \mathfrak{T}_1 \quad Q \Vdash \Pi, \Delta_1, I_1 \ \# \ (\Delta_2, a : A \mid I_2 \vdash c : C) \ /\!\!/ \ \mathfrak{T}_2 \quad \mathfrak{T}_1 \therefore \mathfrak{T}_2}{P \parallel_a Q \Vdash \Pi \ \# \ (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C) \ /\!\!/ \ \mathfrak{T}_1 \parallel \mathfrak{T}_2} \ (\text{Par})$$

$$\frac{P \Rightarrow \Pi \ \# \ (\Delta \mid I \vdash [a : A]) \ /\!\!/ \ \mathfrak{T} \quad \text{princ}(P) = a}{P \Vdash \Pi \ \# \ (\Delta \mid I \vdash a : A) \ /\!\!/ \ \mathfrak{T}} \ (\Vdash\text{R})$$

$$\frac{P \Rightarrow \Pi \ \# \ (\Delta, [a : A] \mid I \vdash c : C) \ /\!\!/ \ \mathfrak{T} \quad \text{princ}(P) = a}{P \Vdash \Pi \ \# \ (\Delta, a : A \mid I \vdash c : C) \ /\!\!/ \ \mathfrak{T}} \ (\Vdash\text{L})$$

Fig. 1. Uniform Process Typing Judgment

Intuitively, it specifies that if traces t_1 and t_2 corresponding to P and Q can be composed as process executions, then each synchronized interleaving of t_1 and t_2 qua process traces can be obtained by synchronizing t_1 and t_2 as constraint traces. In other words, we can interleave the constraints in t_1 and t_2 such that all of the constraints are consistent with the process actions in t_1 and t_2, and recover all valid process interleavings. The key point of this definition is that we can recover all valid process interleavings despite the presence of constraints, i.e., that $(t_1/_¿ \parallel t_2/_¿) \subseteq (t_1 \parallel t_2)/_¿$; we already know by proposition 5.5 that adding constraints cannot produce new process interleavings, i.e., that $(t_1/_¿ \parallel t_2/_¿) \supseteq (t_1 \parallel t_2)/_¿$.

Example 6.1. We illustrate how constraint satisfaction rules out incompatible compositions by attempting to compose the following two processes (respectively, P and Q) in parallel:

case $a \ \{l \Rightarrow b.0; \ \text{wait} \ a; \ \text{close} \ b\ \}_{l \in \text{Bit}} \Vdash \cdot \ \# \ (a : \text{Bits}_1 \mid \cdot \vdash b : \text{Bits}_1) \ /\!\!/ \ \mathfrak{T}_1$

case $a \ \{l \Rightarrow b.l; \ \text{wait} \ a; \ \text{close} \ b\ \}_{l \in \text{Bit}} \Vdash a : \text{Bits}_1 \ \# \ (b : \text{Bits}_1 \mid \cdot \vdash c : \text{IdBitSeq}_1 \ a) \ /\!\!/ \ \mathfrak{T}_2.$

Their composition should not satisfy $P \parallel_a Q \Vdash \cdot \ \# \ (a : \text{Bits}_1 \mid b : \text{Bits}_1 \vdash c : \text{IdBitSeq}_1 \ a) \ /\!\!/ \ \mathfrak{T}$ because an input of 1 on a does not result in the output 1 on c specified by $\text{IdBitSeq}_1 \ a$. The following two traces violate (21) (they have no interleavings) and keep (Par) from being applied:

(?; label 1 on a) :: (!; label 0 on b) :: (?; close on a) :: (!; close on b) :: $\varepsilon \in \mathfrak{T}_1$

(?; label 0 on b) :: ($¿$; label 0 on a) :: (?; label 0 on c) :: $\cdots \in \mathfrak{T}_2$

The rules ($\Vdash$R) and ($\Vdash$L) uniformly identify the next channel a on which a process will communicate, and kick off the second phase that reduces the type of a. This channel a is called the **principal channel** of the process, and is given by $\text{princ}(P) = a$ for

$P \coloneqq \text{close} \ a \mid \text{wait} \ a; \ Q \mid a.k; \ Q \mid \text{case} \ a \ \{l \Rightarrow Q_l\}_{l \in L} \mid \text{send} \ a \ (b.Q_1); \ Q_2 \mid b \leftarrow \text{receive} \ a; \ Q.$

We remark that there is exactly one rule in the first phase for each process forming construct.

6.2 Phase two: focussed type reduction

The second phase is driven by a pair of type-reduction judgments $P \Rightarrow \Pi \ \# \ (\Delta \mid I \vdash [a : A]) \ /\!\!/ \ \mathfrak{T}$ and $P \Rightarrow \Pi \ \# \ (\Delta, [a : A] \mid I \vdash c : C) \ /\!\!/ \ \mathfrak{T}$ that reduce a focussed type $[a : A]$ to weak head normal form, and checks that if the principal channel of P is a, then P communicates on a according to A. The right-focussed judgment is inductively defined by the rules of fig. 2 (weak head normal cases) and fig. 3 (reduction cases); the analogous left-focussed judgment exchanges sending and receiving.

$$\boxed{P \Rightarrow \Pi \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : A]) \mathbin{/\!\!/} \mathfrak{T}} \qquad \text{Focussed reduction of } a : A \text{ on the right}$$

$$\frac{}{\text{close } a \Rightarrow \Pi \ \# \ (\cdot \mathbin{\text{\sc i}} \cdot \vdash [a : \mathbf{1}]) \mathbin{/\!\!/} \{\,(!;\ \text{close on } a) :: \varepsilon\,\}} \ (\Rightarrow\!\mathbf{1}\text{R})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi \mathbin{\text{\sc i}} \mathrm{I}/\pi \vdash a : A_k) \mathbin{/\!\!/} \mathfrak{T} \quad \pi = \text{label } k \text{ on } a}{a.k;\ P \Rightarrow \Pi \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : \oplus_{l \in L \cup \{k\}} A_l]) \mathbin{/\!\!/} (!;\ \text{label } k \text{ on } a) :: \mathfrak{T}} \ (\Rightarrow\!\oplus\text{R})$$

$$\frac{P_l \Vdash \Pi/\pi \ \# \ (\Delta/\pi \mathbin{\text{\sc i}} \mathrm{I}/\pi \vdash a : A_l) \mathbin{/\!\!/} \mathfrak{T}_l \quad \pi = \text{label } l \text{ on } a \quad (\forall l \in L)}{\text{case } a \ \{l \Rightarrow P_l\}_{l \in L \cup L'} \Rightarrow \Pi \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : \&_{l \in L} A_l]) \mathbin{/\!\!/} \bigcup_{l \in L} (?;\ \text{label } k \text{ on } a) :: \mathfrak{T}_l} \ (\Rightarrow\!\&\text{R})$$

$$\frac{\begin{array}{l} P \Vdash (\Pi, \Delta_2, \mathrm{I}_2)/\pi, a : A \ \# \ (\Delta_1/\pi \mathbin{\text{\sc i}} \mathrm{I}_1/\pi \vdash b : B) \mathbin{/\!\!/} \mathfrak{T}_1 \\ Q \Vdash (\Pi, \Delta_1, \mathrm{I}_1)/\pi, b : B \ \# \ (\Delta_2/\pi \mathbin{\text{\sc i}} \mathrm{I}_2/\pi \vdash a : A) \mathbin{/\!\!/} \mathfrak{T}_2 \quad \pi = \text{chan } b \text{ on } a \quad \mathfrak{T}_1 \mathrel{\therefore} \mathfrak{T}_2 \end{array}}{\text{send } a \ (b.P);\ Q \Rightarrow \Pi \ \# \ (\Delta_1, \Delta_2 \mathbin{\text{\sc i}} \mathrm{I}_1, \mathrm{I}_2 \vdash [a : (b : B) \otimes (a : A)]) \mathbin{/\!\!/} (!;\ \text{chan on } a) :: b \ . \ (\mathfrak{T}_1 \parallel \mathfrak{T}_2)} \ (\Rightarrow\!\otimes\text{R})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi, b : B \mathbin{\text{\sc i}} \mathrm{I}/\pi \vdash a : A) \mathbin{/\!\!/} \mathfrak{T} \quad \pi = \text{chan } b \text{ on } a}{b \leftarrow \text{receive } a;\ P \Rightarrow \Pi \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : (b : B) \multimap (a : A)]) \mathbin{/\!\!/} (?;\ \text{chan on } a) :: b \ . \ \mathfrak{T}} \ (\Rightarrow\!\multimap\text{R})$$

Fig. 2. Right-Focussed Process Typing — Focussed Type in Weak Head Normal Form

$$\frac{P \Rightarrow \Pi/\pi \ \# \ (\Delta/\pi \mathbin{\text{\sc i}} \mathrm{I}/\pi \vdash [a : A]) \mathbin{/\!\!/} \mathfrak{T} \quad \pi = \text{close on } c}{P \Rightarrow \Pi, c : \mathbf{1} \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : \text{CASE } c \ \{\text{close} \Rightarrow A\}]) \mathbin{/\!\!/} (\text{¿};\ \text{close on } c) :: \mathfrak{T}} \ (\Rightarrow\!\mathbf{1}[\text{R}])$$

$$\frac{P \Rightarrow \Pi/\pi_l, c : C_l \ \# \ (\Delta/\pi_l \mathbin{\text{\sc i}} \mathrm{I}/\pi_l \vdash [a : A_l]) \mathbin{/\!\!/} \mathfrak{T}_l \quad \pi_l = \text{label } l \text{ on } c \quad (\forall l \in L) \quad L \subseteq L' \subseteq L'' \quad \odot \in \{\oplus, \&\}}{P \Rightarrow \Pi, c : \odot_{l \in L'} C_l \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : \text{CASE } c \ \{l \Rightarrow A_l\}_{l \in L''}]) \mathbin{/\!\!/} \bigcup_{l \in L} (\text{¿};\ \text{label } l \text{ on } c) :: \mathfrak{T}_l} \ (\Rightarrow\!\oplus\&[\text{R}])$$

$$\frac{P \Rightarrow \Pi/\pi, b : B, c : C \ \# \ (\Delta/\pi \mathbin{\text{\sc i}} \mathrm{I}/\pi \vdash [a : A]) \mathbin{/\!\!/} \mathfrak{T} \quad \pi = \text{chan } b \text{ on } c \quad \odot \in \{\otimes, \multimap\}}{P \Rightarrow \Pi, c : (b : B) \odot (c : C) \ \# \ (\Delta \mathbin{\text{\sc i}} \mathrm{I} \vdash [a : \text{CASE } c \ \{\langle b \rangle \Rightarrow A\}]) \mathbin{/\!\!/} (\text{¿};\ \text{chan on } c) :: b \ . \ \mathfrak{T}} \ (\Rightarrow\!\otimes\multimap[\text{R}])$$

Fig. 3. Right-Focussed Process Typing — Reducing the Type in Focus

For example, the rule ($\Rightarrow\!\mathbf{1}$R) of fig. 2 specifies that the process close a communicates according to the weak head normal type $a : \mathbf{1}$, provided it has empty client and internal contexts. The client context must be empty to ensure that we do not accidentally discard any clients. The internal context must be empty to ensure that we use all internal channels in the specification. We must do so, for the current process may have been spawned by a larger process (typed by channels in Π) that observes channels in I, and we must not allow these channels to be discarded. Because no constraints are imposed on the environment, the output constraint set is exactly $[\![\text{close } a]\!]$.

The rule ($\Rightarrow\!\oplus$R) specifies that a process $a.k;\ P$ is a server for a channel $a : \oplus_{l \in L \cup \{k\}} A_l$ provided P is well-typed when $a : A_k$ and the remainder of the specification has been reduced by label k on a. This reduction ensures that all types that depend on a observe the transmitted label k. The resulting constraint set is given by prefixing constraint set $\mathfrak{T}$ for P by the output action $(!, \text{label } k \text{ on } a)$.

Channel transmission send $a\ (b.P);\ Q$ spawns a process P that provides a fresh channel b, sends b over a, and continues as Q. Processes P and Q cannot share any channels. This is ensured by ($\Rightarrow\!\otimes$R): it treats all channels used or provided by P as ambient in Q's typing judgment and vice-versa. The rule also ensures that their specifications impose compatible constraints, analogously to (Par).

Sometimes, the focussed channel is not yet weak head normal. In these cases, we must reduce it to a weak head normal type (see fig. 3). The rule ($\Rightarrow\!\mathbf{1}[\text{R}]$) reduces a focussed channel of type CASE c { close $\Rightarrow A$ }. This reduction is justified so long as a close message is observed on ambient

channel c. This constraint is captured by prefixing the constraint $(\textrm{¿}, \text{close on } c)$ on the constraint set $\mathfrak{T}$ generated by the premise. The rule $(\Rightarrow\oplus\&[R])$ can generate constraints for a subset of labels in a choice to support composition with processes that only send a subset of the permitted labels. The remaining reduction rules are analogous. We remark that the rules of fig. 3 only reduce types that observe ambient channels. A focussed type that observes a used or provided channel cannot be reduced because we cannot reduce types based on future communications. Indeed, the observation required to reduce the type can only be provided by the process being checked, but the process will communicate on the focussed principal channel before it will communicate on the channel being observed. Types that observe local channels are instead immediately reduced by the rules of fig. 2.

Example 6.2. Set $C = \text{CASE } a \{ \text{close} \Rightarrow \mathbf{1} \}$ and $P_{x,y} = (\text{wait } x; \text{ close } y)$. We illustrate the mechanics of our typechecking algorithm by attempting to typecheck $\nu(b : \mathbf{1}) . (P_{a,b} \parallel_b P_{b,c})$ against the specification $\cdot \# (a : \mathbf{1} \mid \cdot \vdash c : C)$. Typechecking succeeds if we can build a derivation of $\nu(b : \mathbf{1}) . (P_{a,b} \parallel_b P_{b,c}) \Vdash \cdot \# (a : \mathbf{1} \mid \cdot \vdash c : C) \mathbin{/\!\!/} \mathfrak{T}$ for some $\mathfrak{T}$. First, we build a candidate derivation by considering only the inputs to the judgment (the red parts of the tree). The only possible bottom rule in our case is (New), and we use proof search to complete the derivation. The focussed nature of our system ensures that search amounts to inversion: in each case, at most one rule can be used to extend the derivation. Eventually each branch either gets stuck or reaches a leaf (an axiom). The following candidate derivation shows that we can successfully build the red portion of the tree:

$$
\cfrac{
  \cfrac{
    \cfrac{
      \cfrac{\overline{\text{close } b \Rightarrow c : \mathbf{1} \, \# \, (\cdot \mid \cdot \vdash [b : \mathbf{1}]) \mathbin{/\!\!/} \mathcal{L}_1}}{\text{close } b \Rightarrow c : \mathbf{1} \, \# \, (\cdot \mid \cdot \vdash b : \mathbf{1}) \mathbin{/\!\!/} \mathcal{L}_1} \; {\scriptstyle(\Vdash\text{R})}
    }{
      \cfrac{P_{a,b} \Rightarrow c : C \, \# \, ([a : \mathbf{1}] \mid \cdot \vdash b : \mathbf{1}) \mathbin{/\!\!/} \mathcal{L}_2}{P_{a,b} \Vdash c : C \, \# \, (a : \mathbf{1} \mid \cdot \vdash b : \mathbf{1}) \mathbin{/\!\!/} \mathcal{L}_2} \; {\scriptstyle(\Vdash\text{L})}
    } \; {\scriptstyle(\Rightarrow\mathbf{1}\text{L})}
    \qquad
    \cfrac{
      \cfrac{
        \cfrac{\overline{\text{close } c \Vdash \cdot \# (\cdot \mid \cdot \vdash [c : \mathbf{1}]) \mathbin{/\!\!/} \mathcal{R}_1}}{\text{close } c \Vdash a : \mathbf{1} \, \# \, (\cdot \mid \cdot \vdash [c : C]) \mathbin{/\!\!/} \mathcal{R}_2} \; {\scriptstyle(\Rightarrow\mathbf{1}[\text{R}])}
      }{\text{close } c \Vdash a : \mathbf{1} \, \# \, (\cdot \mid \cdot \vdash c : C) \mathbin{/\!\!/} \mathcal{R}_2} \; {\scriptstyle(\Vdash\text{R})}
    }{
      \cfrac{P_{b,c} \Rightarrow a : \mathbf{1} \, \# \, ([b : \mathbf{1}] \mid \cdot \vdash c : C) \mathbin{/\!\!/} \mathcal{R}_3}{P_{b,c} \Vdash a : \mathbf{1} \, \# \, (b : \mathbf{1} \mid \cdot \vdash c : C) \mathbin{/\!\!/} \mathcal{R}_3} \; {\scriptstyle(\Vdash\text{L})}
    } \; {\scriptstyle(\Rightarrow\mathbf{1}\text{L})}
  }{P_{a,b} \parallel_b P_{b,c} \Vdash \cdot \# (a : \mathbf{1} \mid b : \mathbf{1} \vdash c : C) \mathbin{/\!\!/} \mathcal{L}_2 \parallel \mathcal{R}_3} \; {\scriptstyle(\text{Par})}
}{\nu(b : \mathbf{1}) . (P_{a,b} \parallel_b P_{b,c}) \Vdash \cdot \# (a : \mathbf{1} \mid \cdot \vdash c : C) \mathbin{/\!\!/} (\mathcal{L}_2 \parallel \mathcal{R}_3) \setminus \{ b \}} \; {\scriptstyle(\text{New})}
$$

Assuming every branch reaches a leaf, the axioms specify the leaf's judgment's output (the trace set in blue). We thread these trace sets back down through the derivation tree, using the operations specified in each rule to build that rules output trace set. In this case:

$$\mathcal{L}_1 = \{ \, (!; \text{ close on } b) :: \varepsilon \, \} \qquad \mathcal{L}_2 = (?; \text{ close on } a) :: \mathcal{L}_1$$

$$\mathcal{R}_1 = \{ \, (!; \text{ close on } c) :: \varepsilon \, \} \qquad \mathcal{R}_2 = (\textrm{¿}; \text{ close on } a) :: \mathcal{R}_1 \qquad \mathcal{R}_3 = (?; \text{ close on } b) :: \mathcal{R}_2$$

$$\mathcal{L}_2 \parallel \mathcal{R}_3 = \{ \, (?; \text{ close on } a) :: (\sigma; \text{ close on } b) :: (!; \text{ close on } c) :: \varepsilon \, \}$$

$$(\mathcal{L}_2 \parallel \mathcal{R}_3) \setminus \{ b \} = \{ \, (?; \text{ close on } a) :: (!; \text{ close on } c) :: \varepsilon \, \}$$

Typechecking succeeds if the resulting derivation is valid (including the side condition (21) for (Par)).

Our rules determine a terminating typechecking algorithm when all choice types are indexed by finite label sets. Indeed, all constraint sets and their traces are finite, so trace operations terminate, and the alternating phases reduce the complexity of processes and their specifications. Most of the typechecking complexity lies in computing operations on traces. We conjecture that they could efficiently be implemented by representing traces as tries or radix trees.

7　SAFETY PROPERTIES

Our typechecking algorithm is sound: if we can typecheck a process against a process specification, then all executions of that process are permitted by its specification. Put differently, our typechecking algorithm ensures that well-typed processes communicate safely:

THEOREM 7.1 (SOUNDNESS). *The typechecking algorithm is sound: if any of*

(1) $P \Vdash \Pi \# (\Delta \mid \mathrm{I} \vdash a : A) \mathbin{/\!\!/} \mathfrak{T}$,

(2) $P \Rightarrow \Pi \# (\Delta \mid I \vdash [a : A]) /\!\!/ \mathfrak{T}$, *or*
(3) $P \Rightarrow \Pi \# (\Delta, [a : A] \mid I \vdash c : C) /\!\!/ \mathfrak{T}$,

then $\mathfrak{T}/_{\dot{\iota}} = [\![P]\!]$ *(the set $\mathfrak{T}$ extents the executions of P with constraints) and* $\mathfrak{T} \subseteq [\![\Pi \# \Delta \mid I \vdash a : A]\!]$
(the constraint set $\mathfrak{T}$ satisfies the specification $\Pi \# \Delta \mid I \vdash a : A$).

COROLLARY 7.2. *If* $P \Vdash \cdot \# (\Delta \mid I \vdash a : A) /\!\!/ \mathfrak{T}$, *then* $\mathfrak{T} = [\![P]\!]$ *and* $[\![P]\!] \subseteq [\![\cdot \# (\Delta \mid I \vdash a : A)]\!]$.

A key feature of MOST is its modular approach to process specification. Processes whose specifications impose no constraints on their environment can always be composed:

PROPOSITION 7.3 (COMPOSITIONALITY). *If* $P \Vdash \cdot \# (\Delta_1 \mid I_1 \vdash a : A) /\!\!/ \mathfrak{T}_1$ *and*
$Q \Vdash \cdot \# (\Delta_2, a : A \mid I_2 \vdash c : C) /\!\!/ \mathfrak{T}_2$, *then* $P \parallel_a Q \Vdash \cdot \# (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C) /\!\!/ \mathfrak{T}_1 \parallel \mathfrak{T}_2$.

8 EXTENDING MOST

MOST provides a foundation for the development of session-typed languages with rich specifications. In this section, we conjecturally sketch extensions along several axes that allow MOST to capture richer computational phenomena: selection, asynchrony, value dependency, and recursion.

8.1 Selection

Go [The Go Project 2024] includes a *select* statement that, given a collection of communication actions to perform, randomly performs one from the set of actions that are possible. For example, if processes P and Q attempt to communicate with a process R, then process R could use a select statement to communicate with the first of the two processes that is ready. We show how to extend MOST with a linear variant of a select statement building on ideas from differential linear logic [Ehrhard 2018] and its applications to concurrency [Rocha and Caires 2021, 2023]. Differential linear logic offers a means of *adding* proofs of the same sequent. In the intuitionistic case, if we ascribe the sequents of intuitionistic linear logic with proof terms (processes), we get the rule

$$\frac{\Gamma \vdash P :: a : A \quad \Gamma \vdash Q :: a : A}{\Gamma \vdash P + Q :: a : A}$$

The possible behaviours of $P + Q$ are all those of P or Q, so the sum denotes $[\![P + Q]\!] = [\![P]\!] \cup [\![Q]\!]$. This clause satisfies the identities described by Rocha and Caires, namely, that process composition distributes over sums and that all processes are idempotent with regard to the sum, i.e., that $P + P \equiv P$ for all P. Our discussion suggests the following typing rule for sums of MOST processes:

$$\frac{P_1 \Vdash \Pi \# (\Delta \mid I \vdash a : A) /\!\!/ \mathcal{T}_1 \quad P_2 \Vdash \Pi \# (\Delta \mid I \vdash a : A) /\!\!/ \mathcal{T}_2}{P_1 + P_2 \Vdash \Pi \# (\Delta \mid I \vdash a : A) /\!\!/ \mathcal{T}_1 \cup \mathcal{T}_2} \text{ (SUM)}$$

Returning to our motivating example, suppose that we want to receive channels a and c over b and d, respectively, but that we do not want to impose a particular order on which of b or d is used first. The following sum of processes receives on whichever of b or d is ready first, non-deterministically breaking any ties, and communicates over the other before continuing as R':

$$(a \leftarrow \text{receive } b; \; c \leftarrow \text{receive } d; \; R') + (c \leftarrow \text{receive } d; \; a \leftarrow \text{receive } b; \; R').$$

8.2 Asynchrony

Adapting MOST to use an asynchronous communication semantics is straightforward, but it comes at the cost of technical complexity. The intended semantics treats channels as unbounded FIFO buffers, where processes can always send (add a message to a buffer), and they can receive a message if one is buffered. We implement FIFO buffers using *continuation channels*, where each channel carries exactly one message, and each message consists of the datum plus the name of the (continuation) channel that will carry the next message. To preclude interference between channels, continuation channel names are globally fresh. To illustrate, consider the trace that describes sending on a the label l,

receiving label k, and closing a. In the synchronous setting, we simply send the three messages on a, giving a trace $(!;\ \text{label}\ l\ \text{on}\ a)::(?;\ \text{label}\ k\ \text{on}\ a)::(!;\ \text{close on}\ a)::\varepsilon$. In the asynchronous setting, the label l is paired with a fresh name a_1 when sent on a; label k is paired with a fresh name a_2 and sent on a_1; and the close message is sent on a_2. To enforce global freshness, ensure that trace sets remain finite, and preserve compositionality, we model continuation channels in traces by bound channel names. The corresponding trace is $(!;\ \text{label}\ l\ \text{on}\ a)::a_1\ .\ (?;\ \text{label}\ k\ \text{on}\ a_1)::a_2\ .\ (!;\ \text{close on}\ a_2)::\varepsilon$, where a_1 and a_2 are bound. Modifying the semantic clauses to capture asynchronous communication is then a matter of adapting the trace operations of section 5 to handle continuation channels, and of threading bindings for continuation channels throughout the semantic clauses. For example, eq. (3) becomes $[\![a.k;\ P]\!] = (!;\ \text{label}\ k\ \text{on}\ a)::a\ .\ [\![P]\!]$, where a is bound in the traces in $[\![P]\!]$. To type processes in the asynchronous setting, we must extend the syntax of observed communications π to include the name of the continuation channel, and extend type reductions A/π to substitute the continuation channel for the carrier of π in A. Because type reduction is localized to the operation A/π and all reductions occur in lockstep in the typing rules, we never need to worry about a continuation channel name going out of scope, and no other changes are required.

8.3 Value-dependent session types

Most can be extended to support value-dependent session types [Toninho, Caires, et al. 2011], where session types can depend on values drawn from an underlying dependent functional language. Value-dependent type formers $\forall x : \tau\ .\ A$ and $\exists x : \tau\ .\ A$ respectively specify servers that receive or send values v of type τ and then communicate according to $[v/x]A$. We first extend our syntax:

$$
\begin{aligned}
\text{Processes} \quad P ::= \cdots \quad &|\quad \text{send}\ M\ \text{on}\ a;\ P && \text{Evaluate term}\ M\ \text{and send its value on}\ a \\
&|\quad \text{let}\ x = \text{rcv}\ a\ \text{in}\ P && \text{Bind received value to}\ x\ \text{in}\ P \\
\text{Session types} \quad A ::= \cdots \quad &|\quad \forall x : \tau\ .\ A && \text{Dependent value receiving} \\
&|\quad \exists x : \tau\ .\ A && \text{Dependent value sending} \\
&|\quad \text{CASE}\ a\ \{\,\text{val}\ x \Rightarrow A\,\} && \text{Value observation}
\end{aligned}
$$

Value observation $\text{CASE}\ a\ \{\,\text{val}\ x \Rightarrow A\,\}$ observes a transmitted value v on a channel a and reduces to the type $[v/x]A$. Its reduction is defined analogously to the other immediate reductions:

$$
(\text{CASE}\ c\ \{\,\text{val}\ x \Rightarrow A\,\})/\pi = \begin{cases} [M/x]A & \text{if}\ \pi = \text{val}\ M\ \text{on}\ c \\ \text{CASE}\ c\ \{\,\text{val}\ x \Rightarrow A/\pi\,\} & \text{if}\ \text{cc}(\pi) \neq c \end{cases}
$$

Write $\mathbf{val}[\tau]$ for the sets of closed values of type τ. We give meaning to value transmission by allowing terms to be embedded in traces:[3]

$$[\![\text{send}\ M\ \text{on}\ a;\ P]\!] = (!;\ \text{val}\ M\ \text{on}\ a)::[\![P]\!], \qquad [\![\text{let}\ x = \text{rcv}\ a\ \text{in}\ P]\!] = \bigcup_{v \in \mathbf{val}[\tau]} (?;\ \text{val}\ v\ \text{on}\ a)::[\![P]\!].$$

We then extend interleaving to use the underlying definitional equality on terms so that whenever M and v are definitionally equal, $((!;\ \text{val}\ M\ \text{on}\ a)::t_1) \| ((?;\ \text{val}\ v\ \text{on}\ a)::t_2) = (\sigma;\ \text{val}\ v\ \text{on}\ a)::(t_1 \| t_2)$

To type processes that use values, we assume a typing judgment $\Gamma \vdash_f M : \tau$ for the functional language and extend the process typing judgment to include Γ. The right process typing rules follow a similar pattern as for label transmission, and the other rules follow analogously:

$$\frac{P \Vdash \Gamma, x : \tau;\ \Pi/\pi \# (\Delta/\pi \mathbin{\mathrm{I}} I/\pi \vdash a : A)\ /\!\!/\ \mathcal{T} \quad \pi = \text{val}\ x\ \text{on}\ a}{\text{let}\ x = \text{rcv}\ a\ \text{in}\ P \Rightarrow \Gamma;\ \Pi \# (\Delta \mathbin{\mathrm{I}} I \vdash [a : \forall x : \tau\ .\ A])\ /\!\!/\ \bigcup_{v \in \mathbf{val}[\tau]} (?;\ \text{val}\ v\ \text{on}\ a)::[v/x]\mathcal{T}}\ (\Rightarrow\forall\text{R})$$

$$\frac{P \Vdash \Gamma;\ \Pi/\pi \# (\Delta/\pi \mathbin{\mathrm{I}} I/\pi \vdash a : [M/x]A)\ /\!\!/\ \mathcal{T} \quad \Gamma \vdash_f M : \tau \quad \pi = \text{val}\ M\ \text{on}\ a}{\text{send}\ M\ \text{on}\ a;\ P \Rightarrow \Gamma;\ \Pi \# (\Delta \mathbin{\mathrm{I}} I \vdash [a : \exists x : \tau\ .\ A])\ /\!\!/\ (!;\ \text{val}\ M\ \text{on}\ a)::\mathcal{T}}\ (\Rightarrow\exists\text{R})$$

---

[3]Extending our semantics to capture evaluation in the underlying dependent type theory is beyond the scope of this sketch.

$$\dfrac{P \Vdash \Gamma, x : \tau; \ \Pi/\pi \ \# \ (\Delta/\pi \mid I/\pi \vdash a : A) \ /\!/ \ \mathcal{T} \quad \pi = \mathsf{val} \ x \ \mathsf{on} \ a \quad \sharp \in \{\forall, \exists\}}{P \Rightarrow \Gamma; \ \Pi, c : (\sharp x : \tau.C) \ \# \ (\Delta \mid I \vdash [a : \mathsf{CASE} \ c \ \{\mathsf{val} \ x \Rightarrow A\}]) \ /\!/ \ \bigcup_{v \in \mathbf{val}[\tau]}(\mathbf{\mathit{i}}; \ \mathsf{val} \ v \ \mathsf{on} \ a) :: [v/x]\mathcal{T}} \ (\Rightarrow\!\forall\exists[\mathrm{R}])$$

8.4   Recursion

Most can be extended to support mutually recursive types and processes using coinductively defined typing judgments. We implement mutually recursive processes using process declarations $f(\vec{x}) = P_f$ gathered in an implicit global context Σ_P. These declarations specify that the process named f is given by its implementation P_f involving channels $\vec{x}$. The syntax call $f(\vec{a})$ calls process f with channel names $\vec{a}$, and it corresponds to running $[\vec{a}/\vec{x}]P_f$. Equirecursive types are given by type declarations $X[\vec{x}] = A_X$ gathered in an implicit global context Σ_A. Channel name variables $\vec{x}$ specify the channels that A_X can observe, and $X[\vec{a}]$ is definitionally equal to $[\vec{a}/\vec{x}]A_X$. A complete program $\mathbb{P}$ is a called process name paired with explicit declarations for each called name.

Processes	$P ::= \cdots \mid \mathsf{call} \ f(\vec{a})$	Call named process f with channels $\vec{a}$
Session types	$A ::= \cdots \mid X[\vec{x}]$	Equirecursive type names
Type declarations	$\Sigma_A ::= \cdot \mid \Sigma_A, X[\vec{x}] = A_X$	$X[\vec{x}]$ is definitionally equal to A_X
Process declarations	$\Sigma_P ::= \cdot \mid \Sigma_P, f(\vec{x}) = P_f$	Process named f is implemented by P_f
Programs	$\mathbb{P} ::= \mathsf{call} \ f(\vec{a}) \ \mathsf{where} \ \Sigma_P$	Complete or top-level program

We assume that recursive declarations are contractive [Pierce 2002, p. 300]. This standard assumption simplifies our theory and allows us to interpret equirecursive types as regular trees. To ensure type reduction A/π remains well defined, we interpret its definition coinductively instead of inductively.

Example 8.1. For each fixed type A, the type $\mathsf{list}_A = \oplus\{\ \mathsf{nil} : \mathbf{1}, \mathsf{cons} : A \otimes \mathsf{list}_A\ \}$ specifies lists of channels of type A. The type $\mathsf{stream}_A = \&\{\ \mathsf{head} : A, \mathsf{tail} : \mathsf{stream}_A\ \}$ specifies a stream of channels of type A. Given a process declaration $\mathsf{u}(x) = (\mathsf{case} \ x \ \{\ \mathsf{head} \Rightarrow \mathsf{close} \ x \mid \mathsf{tail} \Rightarrow \mathsf{call} \ \mathsf{u}(x)\ \})$, the process "$\mathsf{call} \ \mathsf{u}(a)$" is a closed process providing $a : \mathsf{stream}_1$.

Example 8.2. The types Bits and IdBitSeq a and the process id from sections 2.1 and 2.2 are encoded by the following declarations, where we use the label \$ and type $\mathbf{1}$ to specify finite sequences:

$$\mathsf{Bits} = \oplus\{\ \$: \mathbf{1}, \mathtt{0} : \mathsf{Bits}, \mathtt{1} : \mathsf{Bits}\ \}$$

$$\mathsf{IdBitSeq}[x] = \mathsf{CASE} \ x \ \{\ \$ \Rightarrow \oplus\{\ \$: \mathsf{CASE} \ x \ \{\ \mathsf{close} \Rightarrow \mathbf{1}\ \}\ \}$$

$$\mid \mathtt{0} \Rightarrow \oplus\{\ \mathtt{0} : \mathsf{IdBitSeq}[x]\ \} \mid \mathtt{1} \Rightarrow \oplus\{\ \mathtt{1} : \mathsf{IdBitSeq}[x]\ \}\ \}$$

$$\mathsf{id}(x, y) = \mathsf{case} \ x \ \{\ \$ \Rightarrow \mathsf{wait} \ x; \ \mathsf{close} \ y \mid \mathtt{0} \Rightarrow y.\mathtt{0}; \ \mathsf{call} \ \mathsf{id}(x, y) \mid \mathtt{1} \Rightarrow y.\mathtt{1}; \ \mathsf{call} \ \mathsf{id}(x, y)\ \}.$$

We extend our semantics to support calling named processes. Let $\mathcal{P}$ be the set of declared calling interfaces $f(\vec{x})$. A process environment is a map $\rho : \mathcal{P} \to \mathbf{TSets}$ giving the meaning of declared processes, where $\mathbf{TSets}$ is the lattice of sets of traces ordered by inclusion. We generalize $[\![P]\!]$ from an element of $\mathbf{TSets}$ to a continuous map $[\![P]\!] : (\mathcal{P} \to \mathbf{TSets}) \to \mathbf{TSets}$. Calling $f(\vec{x})$ with channels $\vec{a}$ means retrieving its meaning from the environment and instantiating it with the names $\vec{a}$:

$$[\![\mathsf{call} \ f(\vec{a})]\!]\rho = \{\ [\vec{a}/\vec{x}]t \mid t \in \rho(f(\vec{x}))\ \}.$$

We adapt the existing clauses by threading through the environments, e.g.: $[\![P\|_a Q]\!]\rho = [\![P]\!]\rho \| [\![Q]\!]\rho$. To capture parallel composition in the presence of potentially infinite traces, we must extend $\|$ to be a fair merge operator. This ensures that each element from each input traces appears at a finite depth in the interleavings. We omit the details due to space constraints.

Complete programs are processes executed in an environment where every called process name has a corresponding declaration. Their semantics are given by

$$[\![\mathsf{call} \ f(\vec{a}) \ \mathsf{where} \ \Sigma_P]\!] = [\![\mathsf{call} \ f(\vec{a})]\!][\![\Sigma_P]\!], \ \mathsf{where}$$

$$[\![f_1(\vec{x}_1) = P_1, \ldots, f_n(\vec{x}_n) = P_n]\!] = \mathsf{FIX}(\lambda\rho \in (\mathcal{P} \to \mathbf{TSets}). \ [\rho \mid f_1(\vec{x}_1) \mapsto [\![P_1]\!]\rho \mid \cdots \mid f_n(\vec{x}_n) \mapsto [\![P_n]\!]\rho]),$$

FIX is the continuous fixed-point operator on $\mathcal{P} \to \mathbf{TSets}$, and $[g \mid y_1 \mapsto v_1 \mid \cdots \mid y_n \mapsto v_n]\,(x)$ is v_i if $x = y_i$ and $g(x)$ otherwise. Intuitively, $[\![\Sigma_P]\!]$ maps each $f_i(\vec{x}_i)$ to the meaning of $[\![P_i]\!]$ in the fixed point giving meaning to the mutually recursive declarations.

Example 8.3. If $\Sigma_P = (f(x) = x.1;\ \text{call } g(x)),\quad (g(x) = x.0;\ \text{call } f(x))$, then

$$[\![\Sigma_P]\!] = \mathsf{FIX}\,(\lambda \rho.\,[\rho \mid f \mapsto (!;\ \text{label } 1 \text{ on } x) :: \rho(g) \mid g \mapsto (!;\ \text{label } 0 \text{ on } x) :: \rho(f)])$$

$$[\![\text{call } f(a) \text{ where } \Sigma_P]\!] = \{\,(!;\ \text{label } 1 \text{ on } a) :: (!;\ \text{label } 0 \text{ on } a) :: (!;\ \text{label } 1 \text{ on } a) :: \cdots\,\}.$$

Process specifications still denote sets $[\![\Pi \,\#\, \mathcal{G}]\!]$ of traces satisfying the specification. However, to capture recursive behaviour, we adapt its definition to use a coinductively defined elementhood relation. Elementhood is given by a coinductive reading of the rules defining $t \in [\![\Pi \,\#\, \mathcal{G}]\!]$ generated by the obvious analogs of section 4.1, e.g.:

$$\frac{t \in [\![\Pi/\pi \,\#\, \mathcal{G}/\pi]\!] \quad \pi = \text{close on } a}{(!;\ \text{close on } a) :: t \in [\![\Pi \,\#\, \mathcal{G} \mid (\cdot \mid \cdot \vdash a : \mathbf{1})]\!]} \qquad \frac{t \in [\![\Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi \mid \mathrm{I}/\pi \vdash a : A_k)]\!] \quad \pi = \text{label } k \text{ on } a}{(!;\ \text{label } k \text{ on } a) :: t \in [\![\Pi \,\#\, \mathcal{G} \mid (\Delta \mid \mathrm{I} \vdash a : \oplus_{l \in L} A_l)]\!]}$$

In order to type mutually recursive processes, we give our typing rules a coinductive reading and treat each component as an input to the judgment. Then, we extend our rules with:

$$\frac{[\vec{a}/\vec{x}]P_f \Vdash \Pi \,\#\, (\Delta \mid \mathrm{I} \vdash c : C) \mathbin{/\!\!/} \mathcal{T} \quad (f(\vec{x}) = P_f) \in \Sigma_P}{\text{call } f(\vec{a}) \Vdash \Pi \,\#\, (\Delta \mid \mathrm{I} \vdash c : C) \mathbin{/\!\!/} \mathcal{T}} \ (\text{Call})$$

This system can be given a finitary presentation by moving from arbitrary coinductive derivations to circular derivations (finite derivations with loops) [Derakhshan and Pfenning 2022; Fortier and Santocanale 2013]. We refer the reader to [Somayyajula and Pfenning 2022] for a discussion of how to do so to achieve typechecking in finite time, and also a discussion of how to encode inductive and coinductive types using value dependency and (general) equirecursive types.

9 RELATED WORK

Most builds on linear-logical foundations of session types pioneered by Caires and Pfenning [2010]. Indeed, the simply typed fragment of our system is based on a proofs-as-processes interpretation of multiplicative-additive intuitionistic linear logic (MAILL), where propositions correspond to binary session types that specify communications on communication channels. A key difference with our work is that such interpretations typically combine parallel composition and hiding into a single operation. This is because process composition corresponds to the cut rule of linear logic. In contrast, we have separate parallel composition and hiding operators, and we track internal channels. This separation ensures that composition does not hide channels observed by specifications, thereby ensuring that specifications remain well-defined after process composition. We can recover cut-style process composition with a derived rule. Our uniform process typing system can then be seen as a conservative extension of typical interpretations of identity-free MAILL. Another key difference is our use of focussing. Processes can be specified with observing types, but we can only check specific communication actions (e.g., sending a label)) against weak head normal types (e.g., an internal choice). This means that to type check a process, we must reduce the type of its principal channel to weak head normal form, potentially generating constraints on ambient communications along the way. Our focussing approach ensures that we only reduce the types of principal channels. This avoids the premature or spurious generation of constraints on ambient channels that could affect constraint compatibility in process composition, and it ensures that type checking is deterministic.

There are various dependent type systems for binary session types. Toninho, Caires, et al. [2011] and Toninho and Yoshida [2018] interpret logical quantifiers as value-dependent session types in languages with functional programming features. As discussed in section 8.3, value-dependent session types depend on transmitted functional values. Variations include label-dependent session

types [Thiemann and Vasconcelos 2019], which reify labels as a first-class type, and arithmetic refinements [Das and Pfenning 2020], which offer type-level arithmetic using transmitted natural numbers. Like other binary session types, value-dependent session types take a channel-local approach: a type only considers communication on a single channel. In contrast, types in Most specify protocols on individual channels while potentially depending on ambient communications.

Multiparty session types [Honda, Yoshida, et al. 2016] are an alternative approach to specifying systems of communicating processes. Instead of specifying communications on individual channels, multiparty session types specify interactions between multiple parties using a global type, and then for each party project a local type specifying its interactions with the other parties. Though multiparty session types can specify complex interactions, they are not compositional and typically require that entire systems be specified at once. To achieve compositionality, Stolze et al. [2023] introduced *partial multiparty sessions types*, which specify multiparty sessions where some parties may be missing. Determining when partial sessions can be composed and the resulting composition is technically complex. In contrast, Most is designed to ensure compositional process specifications, and the composition of process specifications is determined by a handful of inference rules.

The dynamics of processes in proofs-as-processes interpretations of intuitionistic linear logic is often given by multiset rewriting semantics [Cervesato and Scedrov 2009]. We instead describe the behaviour of processes using a trace semantics. We do so to simplify relating the meaning of processes (a set of traces) with the meaning of process specifications (a set of allowed traces). It also provides a means of ensuring that processes and specifications are mutually compatible when typechecking their composition. We conjecture that our trace semantics captures the same observable communication behaviours [Atkey 2017; Kavanagh 2022] as multiset-based semantics. It is challenging to define trace semantics or ordered semantics for process calculi with name generation or α-conversion. Some approaches, e.g., typed event structures for π-calculi [Varacca and Yoshida 2010], specify in advance all names that will be created; others use complex semantic composition operators involving renamings [Crafa et al. 2007]. In contrast, we use binding to range over all possible fresh names in a manner reminiscent of higher-order abstract syntax. This avoids unintended name clashes when interleaving traces during process composition. Our compact presentation also minimizes the number of traces considered by our typechecking algorithm. Traces with binding are similar to nominal sequences [Gabbay and Ghica 2012], which use a coabstraction operator to bind atoms appearing later in a sequence. In contrast, our traces bind not atoms, but names found in later messages. As a result, our traces more closely resemble abstract binding trees.

10   CONCLUSION

In this work we describe Most, a message-observing binary session type system. It uses type-level processes to specify protocols that vary based on ambient communications. This allows us to express more precise safety guarantees that cannot be expressed by prior work.

Our main goal in designing Most is to ensure compositionality. We achieve this using a novel trace-based semantics that uses traces with bindings to compactly handle higher-order communications. Our main technical results are type safety and compositionality for well-typed processes.

In the future, we plan to extend Most with shared channels. This would allow Most to capture shared data structures like shared queues, and shared services like databases.

ACKNOWLEDGMENTS

A OMITTED RULES

The following rules are those of figs. 2 and 3, plus any rules omitted from those figures:

$$\boxed{\begin{array}{l} P \Rightarrow \Pi \ \# \ (\Delta \mid I \vdash [a : A]) \ /\!\!/ \ \mathfrak{T} \\ P \Rightarrow \Pi \ \# \ (\Delta, [a : A] \mid I \vdash c : C) \ /\!\!/ \ \mathfrak{T} \end{array}} \qquad \text{Focussed reduction of } a : A$$

$$\frac{}{\text{close } a \Rightarrow \Pi \ \# \ (\cdot \mid \cdot \vdash [a : 1]) \ /\!\!/ \ \{ \ (!; \ \text{close on } a) :: \varepsilon \ \}} \ (\Rightarrow 1\text{R})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi \mid I/\pi \vdash c : C/\pi) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{close on } a}{\text{wait } a; \ P \Rightarrow \Pi \ \# \ (\Delta, [a : 1] \mid I \vdash c : C) \ /\!\!/ \ (?; \ \text{close on } a) :: \mathfrak{T}} \ (\Rightarrow 1\text{L})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi \mid I/\pi \vdash a : A_k) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{label } k \text{ on } a}{a.k; \ P \Rightarrow \Pi \ \# \ (\Delta \mid I \vdash [a : \oplus_{l \in L \cup \{ k \}} A_l]) \ /\!\!/ \ (!; \ \text{label } k \text{ on } a) :: \mathfrak{T}} \ (\Rightarrow \oplus\text{R})$$

$$\frac{P_l \Vdash \Pi/\pi \ \# \ (\Delta/\pi, a : A_l \mid I/\pi \vdash c : C/\pi) \ /\!\!/ \ \mathfrak{T}_l \quad \pi = \text{label } l \text{ on } a \quad (\forall l \in L)}{\text{case } a \ \{ l \Rightarrow P_l \}_{l \in L \cup L'} \Rightarrow \Pi \ \# \ (\Delta, [a : \oplus_{l \in L} A_l] \mid I \vdash c : C) \ /\!\!/ \ \bigcup_{l \in L}(?; \ \text{label } k \text{ on } a) :: \mathfrak{T}_l} \ (\Rightarrow \oplus\text{L})$$

$$\frac{P_l \Vdash \Pi/\pi \ \# \ (\Delta/\pi \mid I/\pi \vdash a : A_l) \ /\!\!/ \ \mathfrak{T}_l \quad \pi = \text{label } l \text{ on } a \quad (\forall l \in L)}{\text{case } a \ \{ l \Rightarrow P_l \}_{l \in L \cup L'} \Rightarrow \Pi \ \# \ (\Delta \mid I \vdash [a : \&_{l \in L} A_l]) \ /\!\!/ \ \bigcup_{l \in L}(?; \ \text{label } k \text{ on } a) :: \mathfrak{T}_l} \ (\Rightarrow \&\text{R})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi, a : A_k \mid I/\pi \vdash c : C/\pi) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{label } k \text{ on } a}{a.k; \ P \Rightarrow \Pi \ \# \ (\Delta, [a : \&_{l \in L \cup \{ k \}} A_l] \mid I \vdash c : C) \ /\!\!/ \ (!; \ \text{label } k \text{ on } a) :: \mathfrak{T}} \ (\Rightarrow \&\text{L})$$

$$\frac{\begin{array}{c} P \Vdash (\Pi, \Delta_2, I_2)/\pi, a : A \ \# \ (\Delta_1/\pi \mid I_1/\pi \vdash b : B) \ /\!\!/ \ \mathfrak{T}_1 \\ Q \Vdash (\Pi, \Delta_1, I_1)/\pi, b : B \ \# \ (\Delta_2/\pi \mid I_2/\pi \vdash a : A) \ /\!\!/ \ \mathfrak{T}_2 \quad \pi = \text{chan } b \text{ on } a \quad \mathfrak{T}_1 \therefore \mathfrak{T}_2 \end{array}}{\text{send } a \ (b.P); \ Q \Rightarrow \Pi \ \# \ (\Delta_1, \Delta_2 \mid I_1, I_2 \vdash [a : (b : B) \otimes (a : A)]) \ /\!\!/ \ (!; \ \text{chan on } a) :: b \ . \ (\mathfrak{T}_1 \parallel \mathfrak{T}_2)} \ (\Rightarrow \otimes\text{R})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi, b : B, a : A \mid I/\pi \vdash c : C/\pi) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{chan } b \text{ on } a}{b \leftarrow \text{receive } a; \ P \Rightarrow \Pi \ \# \ (\Delta, [a : (b : B) \otimes (a : A)] \mid I \vdash c : C) \ /\!\!/ \ (?; \ \text{chan on } a) :: b \ . \ \mathfrak{T}} \ (\Rightarrow \otimes\text{L})$$

$$\frac{P \Vdash \Pi/\pi \ \# \ (\Delta/\pi, b : B \mid I/\pi \vdash a : A) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{chan } b \text{ on } a}{b \leftarrow \text{receive } a; \ P \Rightarrow \Pi \ \# \ (\Delta \mid I \vdash [a : (b : B) \multimap (a : A)]) \ /\!\!/ \ (?; \ \text{chan on } a) :: b \ . \ \mathfrak{T}} \ (\Rightarrow \multimap\text{R})$$

$$\frac{\begin{array}{c} P \Vdash (\Pi, \Delta_2, I_2)/\pi, a : A \ \# \ (\Delta_1/\pi \mid I_1/\pi \vdash b : B) \ /\!\!/ \ \mathfrak{T}_1 \\ Q \Vdash (\Pi, \Delta_1, I_1)/\pi, b : B \ \# \ (\Delta_2/\pi, a : A \mid I_2/\pi \vdash c : C/\pi) \ /\!\!/ \ \mathfrak{T}_2 \quad \pi = \text{chan } b \text{ on } a \quad \mathfrak{T}_1 \therefore \mathfrak{T}_2 \end{array}}{\text{send } a \ (b.P); \ Q \Rightarrow \Pi \ \# \ (\Delta_1, \Delta_2, [a : (b : B) \multimap (a : A)] \mid I \vdash c : C) \ /\!\!/ \ (!; \ \text{chan on } a) :: b \ . \ (\mathfrak{T}_1 \parallel \mathfrak{T}_2)} \ (\Rightarrow \multimap\text{L})$$

$$\frac{P \Rightarrow \Pi/\pi \ \# \ (\Delta/\pi \mid I/\pi \vdash [a : A]) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{close on } c}{P \Rightarrow \Pi, c : 1 \ \# \ (\Delta \mid I \vdash [a : \text{CASE } c \ \{ \text{close} \Rightarrow A \}]) \ /\!\!/ \ (\text{¿}; \ \text{close on } c) :: \mathfrak{T}} \ (\Rightarrow 1[\text{R}])$$

$$\frac{P \Rightarrow \Pi/\pi \ \# \ (\Delta/\pi, [a : A] \mid I/\pi \vdash d : D/\pi) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{close on } c}{P \Rightarrow \Pi, c : 1 \ \# \ (\Delta, [a : \text{CASE } c \ \{ \text{close} \Rightarrow A \}] \mid I \vdash d : D) \ /\!\!/ \ (\text{¿}; \ \text{close on } c) :: \mathfrak{T}} \ (\Rightarrow 1[\text{L}])$$

$$\frac{P \Rightarrow \Pi/\pi_l, c : C_l \ \# \ (\Delta/\pi_l \mid I/\pi_l \vdash [a : A_l]) \ /\!\!/ \ \mathfrak{T}_l \quad \pi_l = \text{label } l \text{ on } c \quad (\forall l \in L) \quad L \subseteq L' \subseteq L'' \quad \odot \in \{ \oplus, \& \}}{P \Rightarrow \Pi, c : \odot_{l \in L'} C_l \ \# \ (\Delta \mid I \vdash [a : \text{CASE } c \ \{ l \Rightarrow A_l \}_{l \in L''}]) \ /\!\!/ \ \bigcup_{l \in L}(\text{¿}; \ \text{label } l \text{ on } c) :: \mathfrak{T}_l} \ (\Rightarrow \oplus\&[\text{R}])$$

$$\frac{P \Rightarrow \Pi/\pi_l, c : C_l \ \# \ (\Delta/\pi_l, [a : A_l] \mid I/\pi_l \vdash d : D/\pi_l) \ /\!\!/ \ \mathfrak{T}_l \quad \pi_l = \text{label } l \text{ on } c \quad (\forall l \in L) \quad \odot \in \{ \oplus, \& \}}{P \Rightarrow \Pi, c : \odot_{l \in L} C_l \ \# \ (\Delta, [a : \text{CASE } c \ \{ l \Rightarrow A_l \}_{l \in L \cup L'}] \mid I \vdash d : D) \ /\!\!/ \ \bigcup_{l \in L}(\text{¿}; \ \text{label } l \text{ on } c) :: \mathfrak{T}_l} \ (\Rightarrow \oplus\&[\text{L}])$$

$$\frac{P \Rightarrow \Pi/\pi, b : B, c : C \ \# \ (\Delta/\pi \mid I/\pi \vdash [a : A]) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{chan } b \text{ on } c \quad \odot \in \{ \otimes, \multimap \}}{P \Rightarrow \Pi, c : (b : B) \odot (c : C) \ \# \ (\Delta \mid I \vdash [a : \text{CASE } c \ \{ \langle b \rangle \Rightarrow A \}]) \ /\!\!/ \ (\text{¿}; \ \text{chan on } c) :: b \ . \ \mathfrak{T}} \ (\Rightarrow \otimes\multimap[\text{R}])$$

$$\frac{P \Rightarrow \Pi/\pi, b : B, c : C \ \# \ (\Delta/\pi, [a : A] \mid I/\pi \vdash d : D/\pi) \ /\!\!/ \ \mathfrak{T} \quad \pi = \text{chan } b \text{ on } c \quad \odot \in \{ \otimes, \multimap \}}{P \Rightarrow \Pi, c : (b : B) \odot (c : C) \ \# \ (\Delta, [a : \text{CASE } c \ \{ \langle b \rangle \Rightarrow A \}] \mid I \vdash d : D) \ /\!\!/ \ (\text{¿}; \ \text{chan on } c) :: b \ . \ \mathfrak{T}} \ (\Rightarrow \otimes\multimap[\text{L}])$$

B OMITTED SEMANTIC CLAUSES

The following clauses were omitted from section 4.1.

Where $\pi = $ label k on a, $m = $ label k on a, and $k \in L$:

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta \mid \mathrm{I} \vdash a : \&_{l \in L} A_l) \rrbracket \supseteq (?;\ m) :: \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi \mid \mathrm{I}/\pi \vdash a : A_k) \rrbracket \tag{22}$$

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta, a : \&_{l \in L} A_l \mid \mathrm{I} \vdash c : C) \rrbracket \supseteq (!;\ m) :: \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi, a : A_k \mid \mathrm{I}/\pi \vdash c : C/\pi) \rrbracket \tag{23}$$

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta \mid \mathrm{I}, a : \&_{l \in L} A_l \vdash c : C) \rrbracket \supseteq (\sigma;\ m) :: \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi \mid \mathrm{I}/\pi, a : A_k \vdash c : C/\pi) \rrbracket \tag{24}$$

Where $\pi = $ (chan b on a) with b fresh, and $m = $ chan on a:

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta \mid \mathrm{I}, a : (b : B) \otimes (a : A) \vdash c : C) \rrbracket$$
$$\supseteq (\sigma;\ \pi) :: b \,.\, \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi \mid \mathrm{I}/\pi, b : B, a : A \vdash c : C/\pi) \rrbracket \tag{25}$$

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta \mid \mathrm{I} \vdash a : (b : B) \multimap (a : A)) \rrbracket$$
$$\supseteq (?;\ m) :: b \,.\, \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi, b : B \mid \mathrm{I}/\pi \vdash a : A) \rrbracket \tag{26}$$

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta \mid \mathrm{I}, a : (b : B) \multimap (a : A) \vdash c : C) \rrbracket$$
$$\supseteq (\sigma;\ \pi) :: b \,.\, \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta/\pi \mid \mathrm{I}/\pi, b : B, a : A \vdash c : C/\pi) \rrbracket \tag{27}$$

$$\llbracket \Pi \,\#\, \mathcal{G} \mid (\Delta_1, \Delta_2, a : (b : B) \multimap (a : A) \mid \mathrm{I}_1, \mathrm{I}_2 \vdash c : C) \rrbracket$$
$$\supseteq (!;\ m) :: b \,.\, \llbracket \Pi/\pi \,\#\, \mathcal{G}/\pi \mid (\Delta_1/\pi \mid \mathrm{I}_1/\pi \vdash b : B) \mid (\Delta_2/\pi, a : A \mid \mathrm{I}_2/\pi \vdash c : C/\pi) \rrbracket \tag{28}$$

C PROOFS

Proposition 3.3 (Termination). *For all processes P and traces $t \in \llbracket P \rrbracket$, t is finite.*

Proof of Proposition 3.3. By induction on the structure of P, observing that each semantic clause extends a finite trace by a finite number of actions. □

To help characterize the behaviour of reduction on traces with the behaviour of type reduction in proposition C.4, we introduce a reduction relation on contexts and specifications, and a few technical lemmas (lemmas C.1 and C.2 and proposition C.3). We assume that bound names are unique and disjoint from free names. The relation is defined by the following schemata, which hold only if both sides of the arrow are defined:

$$\Pi, a : \mathbf{1} \xRightarrow{\text{close on } a} \Pi/(\text{close on } a) \tag{29}$$

$$\Pi, a : \odot_{l \in L} A_l \xRightarrow{\text{label } l \text{ on } a} \Pi/(\text{label } l \text{ on } a), a : A_l \qquad (l \in L, \odot \in \{ \oplus, \& \}) \tag{30}$$

$$\Pi, a : (b : B) \odot (a : A) \xRightarrow{\text{chan } b \text{ on } a} \Pi/(\text{chan } b \text{ on } a), b : B, a : A \qquad (\odot \in \{ \otimes, \multimap \}) \tag{31}$$

The relation $\xRightarrow{(¿,\pi)}$ on pairs of well-formed specifications is defined as:

$$\Pi \,\#\, \mathcal{G} \xRightarrow{(¿,\pi)} \Pi' \,\#\, \mathcal{G}/\pi \qquad \text{if } \Pi \xRightarrow{\pi} \Pi' \tag{32}$$

We invite the reader to compare reduction (32) with the clauses generating traces of the form $(¿;\ m) :: \vec{x} \,.\, t$ in section 4.1.

Lemma C.1 (Partial Commutativity of Reduction). *Let π_1 and π_2 be observed communications such that $\mathrm{fc}(\pi_1)$ and $\mathrm{fc}(\pi_2)$ are disjoint.*

(1) Both A/π_1 and A/π_2 are defined if and only if both $(A/\pi_1)/\pi_2$ and $(A/\pi_2)/\pi_1$ are defined.
(2) If $(A/\pi_1)/\pi_2$ and $(A/\pi_2)/\pi_1$ are both defined, then they are equal.

PROOF OF LEMMA C.1. It is immediate that if $(A/\pi_1)/\pi_2$ and $(A/\pi_2)/\pi_1$ are both defined, then so are A/π_1 and A/π_2. Assume that A/π_1 and A/π_2 are both defined. We show by induction on the structure of A that $(A/\pi_1)/\pi_2$ and $(A/\pi_2)/\pi_1$ are both defined and equal.

The weak-head normal cases follow readily by the induction hypotheses.

Assume next that $A = \text{CASE } c \{ \text{close} \Rightarrow B \}$. We proceed by case analysis on π_1. Assume first that $\pi_1 = \text{close on } c$. In this case, $\text{cc}(\pi_2) \neq c$, so $A/\pi_2 = \text{CASE } c \{ \text{close} \Rightarrow B/\pi_2 \}$ is defined, so B/π_2 is defined. Then $(A/\pi_1)/\pi_2 = B/\pi_2 = (A/\pi_2)/\pi_1$ are both defined and equal. Otherwise, if neither π_1 or π_2 is close on c, then $\text{cc}(\pi_1) \neq c$ and $\text{cc}(\pi_2) \neq c$. It follows that B/π_i is defined for $i = 1, 2$. We then deduce by the induction hypothesis:

$$(A/\pi_1)/\pi_2 = \text{CASE } c \{ \text{close} \Rightarrow (B/\pi_1)/\pi_2 \} = \text{CASE } c \{ \text{close} \Rightarrow (B/\pi_2)/\pi_1 \} = (A/\pi_2)/\pi_1.$$

Assume next that $A = \text{CASE } c \{ l \Rightarrow A_l \}_{l \in L}$. We proceed by case analysis on π_1. Assume first that $\pi_1 = \text{label } l \text{ on } c$, where $l \in L$. In this case, $(A/\pi_1)/\pi_2 = A_l/\pi_2 = (A/\pi_2)/\pi_1$ are defined and equal. Otherwise, $\text{cc}(\pi_1) \neq c$ and $\text{cc}(\pi_2) \neq c$. In this case, A_l/π_i is defined for $l \in L$ and $i = 1, 2$. We deduce by the induction hypothesis that the following are defined and equal:

$$(A/\pi_1)/\pi_2 = \text{CASE } c \{ l \Rightarrow (A_l/\pi_1)/\pi_2 \}_{l \in L} = \text{CASE } c \{ l \Rightarrow (A_l/\pi_2)/\pi_1 \}_{l \in L} = (A/\pi_2)/\pi_1.$$

The case $A = \text{CASE } c \{ \langle b \rangle \Rightarrow B \}$ is analogous.   □

The following lemma is used to relate two different reductions in the proof of the diamond property (proposition C.3):

LEMMA C.2. *If Π/π_1 is defined, $\Pi \overset{\pi_2}{\Longrightarrow} \Gamma$, and $\text{fc}(\pi_1) \cap \text{fc}(\pi_2) = \emptyset$, then $\Pi/\pi_1 \overset{\pi_2}{\Longrightarrow} \Gamma/\pi_1$.*

PROOF OF LEMMA C.2. By case analysis on the transition $\Pi \overset{\pi_2}{\Longrightarrow} \Gamma$.

CASE (29). Assume the transition is given by $\Pi = \Delta, a : \mathbf{1} \overset{\pi_2}{\Longrightarrow} \Delta/\pi_2 = \Gamma$ where $\pi_2 = \text{close on } a$. Then Δ/π_1 and Δ/π_2 are both defined, so $(\Delta/\pi_1)/\pi_2 = (\Delta/\pi_2)/\pi_1$ are both defined by lemma C.1. It follows from (29) that:

$$\Pi/\pi_1 = \Delta/\pi_1, a : \mathbf{1} \overset{\pi_2}{\Longrightarrow} (\Delta/\pi_1)/\pi_2 = (\Delta/\pi_2)/\pi_1 = \Gamma/\pi_1.$$

CASE (30). Assume the transition is given by $\Pi = \Delta, a : \odot_{l \in L} A_l \overset{\pi_2}{\Longrightarrow} \Delta/\pi_2, a : A_l = \Gamma$ with $\pi_2 = \text{label } l \text{ on } a$, and $\odot \in \{ \oplus, \& \}$. Then Δ/π_1 and Δ/π_2 are both defined, so $(\Delta/\pi_1)/\pi_2 = (\Delta/\pi_2)/\pi_1$ are both defined by lemma C.1. It follows from (30) that:

$$\Pi/\pi_1 = \Delta/\pi_1, a : \odot_{l \in L} A_l/\pi_1 \overset{\pi_2}{\Longrightarrow} (\Delta/\pi_1)/\pi_2, a : A_l/\pi_1 = (\Delta/\pi_2, a : A_l)/\pi_1 = \Gamma/\pi_1.$$

CASE (31). Assume the transition is given by $\Pi = \Delta, a : (b : B) \odot (a : A) \overset{\pi_2}{\Longrightarrow} \Delta/\pi_2, b : B, a : A = \Gamma$ with $\pi_2 = \text{chan } b \text{ on } a$, and $\odot \in \{ \otimes, \multimap \}$. Then Δ/π_1 and Δ/π_2 are both defined, so $(\Delta/\pi_1)/\pi_2 = (\Delta/\pi_2)/\pi_1$ are both defined by lemma C.1. It follows from (31) that:

$$\Pi/\pi_1 = \Delta/\pi_1, a : (b : B/\pi_1) \odot (a : A/\pi_1)$$
$$\overset{\pi_2}{\Longrightarrow} (\Delta/\pi_1)/\pi_2, b : B/\pi_1, a : A/\pi_1 = (\Delta/\pi_2, b : B, a : A)/\pi_1 = \Gamma/\pi_1. \quad □$$

PROPOSITION C.3 (DIAMOND PROPERTY). *If $\Pi \overset{\pi_1}{\Longrightarrow} \Pi_1$ and $\Pi \overset{\pi_2}{\Longrightarrow} \Pi_2$ where $\text{fc}(\pi_1) \cap \text{fc}(\pi_2) = \emptyset$, then there exists a Π_3 such that $\Pi_1 \overset{\pi_2}{\Longrightarrow} \Pi_3$ and $\Pi_2 \overset{\pi_1}{\Longrightarrow} \Pi_3$.*

PROOF OF PROPOSITION C.3. If $\Pi \overset{\pi_1}{\Longrightarrow} \Pi_1$ and $\Pi \overset{\pi_2}{\Longrightarrow} \Pi_2$ where $\mathrm{fc}(\pi_1) \cap \mathrm{fc}(\pi_2) = \emptyset$, then $\Pi = \Delta, a_1 : A_1, a_2 : A_2$ for some Δ and $a_i : A_i$ where $a_i : A_i \overset{\pi_i}{\Longrightarrow} \Delta_i$ for $i = 1, 2$. Then

$$\Pi_1 = (\Delta, a_2 : A_2)/\pi_1, \Delta_1,$$
$$\Pi_2 = (\Delta, a_1 : A_1)/\pi_2, \Delta_2.$$

By lemma C.2,

$$(\Delta, a_2 : A_2)/\pi_1 \overset{\pi_2}{\Longrightarrow} (\Delta/\pi_1)/\pi_2, \Delta_2/\pi_1$$

$$(\Delta, a_1 : A_1)/\pi_2 \overset{\pi_1}{\Longrightarrow} (\Delta/\pi_2)/\pi_1, \Delta_1/\pi_2.$$

In particular, Δ_1/π_2 and Δ_2/π_1 are defined. This implies that the following transitions are defined:

$$\Pi_1 \overset{\pi_2}{\Longrightarrow} (\Delta/\pi_1)/\pi_2, \Delta_2/\pi_1, \Delta_1/\pi_2$$

$$\Pi_2 \overset{\pi_1}{\Longrightarrow} (\Delta/\pi_2)/\pi_1, \Delta_2/\pi_1, \Delta_1/\pi_2.$$

The two right sides are equal by lemma C.1, and we take them to be Π_3. □

The following proposition relates trace reductions and type reductions:

PROPOSITION C.4. *If* $t \in [\![\Pi \# \mathcal{G}]\!]$, $\Pi \# \mathcal{G} \overset{(\iota, \pi)}{\Longrightarrow} \Pi' \# \mathcal{G}/\pi$, *and* t/π *is defined,*
then $t/\pi \in [\![\Pi' \# \mathcal{G}/\pi]\!]$.

PROOF OF PROPOSITION C.4. By induction on the derivation $t \in [\![\Pi \# \mathcal{G}]\!]$.
The base case occurs when $t = \varepsilon$, in which case the result is immediate.
The inductive steps occur when $t = (s; \, m) :: \vec{x} \, . \, t'$ for some $s, m, \vec{x}$ and t'. Each inductive case corresponds to a different possible choice of (s, m). We give a few illustrative cases.
Assume first that $s = \, \text{¿}$ and $\pi = (m \propto \vec{x})$.[4] By inversion, $t \in [\![\Pi \# \mathcal{G}]\!]$ because $t' \in [\![\Pi' \# \mathcal{G}/\pi]\!]$. Observe that $t/\pi = t'$. This gives the result.
In all other cases, the fact that t/π is defined implies that $\mathrm{cc}(\pi) \neq \mathrm{cc}(m)$. This in turn implies that t'/π is defined. Moreover, inversion on $t \in [\![\Pi \# \mathcal{G}]\!]$ implies that $t' \in [\![\Pi_0 \# \mathcal{G}_0]\!]$ for some Π_0 and $\mathcal{G}_0$ such that $\Pi, \overline{\mathcal{G}} \overset{(m \propto \vec{x})}{\Longrightarrow} \Pi_0, \overline{\mathcal{G}_0}$, where we write $\overline{\mathcal{G}}$ for the union of the contexts appearing in $\mathcal{G}$. By the diamond property (proposition C.3), it follows that there exists a Π_0' such that $\Pi_0 \# \mathcal{G}_0 \overset{(\iota, \pi)}{\Longrightarrow} \Pi_0' \# \mathcal{G}_0/\pi$, and such that the following two reductions are defined:

$$\Pi, \overline{\mathcal{G}} \overset{(m \propto \vec{x})}{\Longrightarrow} \Pi_0, \overline{\mathcal{G}_0} \overset{\pi}{\Rightarrow} \Pi_0', \overline{\mathcal{G}_0}/\pi$$

$$\Pi, \overline{\mathcal{G}} \overset{\pi}{\Rightarrow} \Pi', \overline{\mathcal{G}}/\pi \overset{(m \propto \vec{x})}{\Longrightarrow} \Pi_0', \overline{\mathcal{G}_0}/\pi.$$

Observe that $\Pi_0, \overline{\mathcal{G}_0} \overset{\pi}{\Rightarrow} \Pi_0', \overline{\mathcal{G}_0}/\pi$ induces a reduction $\Pi_0 \# \mathcal{G}_0 \overset{(\iota, \pi)}{\Longrightarrow} \Pi_0' \# \mathcal{G}_0/\pi$. By the induction hypothesis on $t' \in [\![\Pi_0 \# \mathcal{G}_0]\!]$ and this reduction, we deduce $t'/\pi \in [\![\Pi_0' \# \mathcal{G}_0/\pi]\!]$. Case analysis on s and m with the reduction $\Pi', \overline{\mathcal{G}}/\pi \overset{(m \propto \vec{x})}{\Longrightarrow} \Pi_0', \overline{\mathcal{G}_0}/\pi$ implies $((s; \, m) :: \vec{x} \, . \, (t'/\pi)) \in [\![\Pi' \# \mathcal{G}/\pi]\!]$. But $t/\pi = ((s; \, p) :: \vec{x} \, . \, (\tau/\pi))$, so this gives the result. □

---

[4]Recall that $(m \propto \vec{x})$ is defined on page 17.

Recall from page 18 that a trace is safely constrained if any carrier channel associated with a constraint sign is not also associated with input, output, or synchronization.

LEMMA C.5. *If* $(o;\ m) :: \vec{x}.\ t_1$ *and* $(¿;\ m) :: \vec{x}.\ t_2$ *are safely constrained, then*

$$(o;\ m) :: \vec{x}.\ ((t_1/¿) \parallel (t_2/¿)) \subseteq (((o;\ m) :: \vec{x}.\ t_1)\ /¿) \parallel (((¿;\ m) :: \vec{x}.\ t_2)\ /¿)\ .$$

PROOF OF LEMMA C.5. We remark that the hypothesis that both traces are safely constrained is essential. Indeed, the following traces are not safely constrained (they send/receive on a while also imposing a constraint on a)

$$(!;\ \text{label } l \text{ on } a) :: (¿;\ \text{close on } a) :: \varepsilon$$
$$(¿;\ \text{label } l \text{ on } a) :: (?;\ \text{close on } a) :: \varepsilon,$$

and they do not satisfy the desired identity:

$(!;\ \text{label } l \text{ on } a) :: (((¿;\ \text{close on } a) :: \varepsilon)\ /¿) \parallel (((?;\ \text{close on } a) :: \varepsilon)\ /¿)$

$= (!;\ \text{label } l \text{ on } a) :: (\varepsilon \parallel ((?;\ \text{close on } a) :: \varepsilon))$

$= \{\ (!;\ \text{label } l \text{ on } a) :: (?;\ \text{close on } a) :: \varepsilon\ \}$

$\not\subseteq \emptyset$

$= ((!;\ \text{label } l \text{ on } a) :: \varepsilon) \parallel ((?;\ \text{close on } a) :: \varepsilon)$

$= (((!;\ \text{label } l \text{ on } a) :: (¿;\ \text{close on } a) :: \varepsilon)\ /¿) \parallel (((¿;\ \text{label } l \text{ on } a) :: (?;\ \text{close on } a) :: \varepsilon)\ /¿)\ .$

Assume that both traces are safely constrained. We examine the traces on the right side of the inclusion and observe:

$$((o;\ m) :: \vec{x}.\ t_1)\ /¿ = (o;\ m) :: \vec{x}.\ (t_1/¿)$$
$$((¿;\ m) :: \vec{x}.\ t_2)\ /¿ = t_2/¿.$$

The hypothesis that the traces are safely constrained implies that $\mathrm{cc}(m)$ does not appear free in $(¿;\ m) :: \vec{x}.\ t_2$. It follows that the interleaving on the right side of the inclusion, which by the last two equalities is $((o;\ m) :: \vec{x}.\ (t_1/¿)) \parallel (t_2/¿)$, is either undefined or given by either a base case (when $t_2/¿ = \varepsilon$) or a commuting case.

If it is given by a base case, then the result follows easily:

$(o;\ m) :: \vec{x}.\ ((t_1/¿) \parallel (t_2/¿))$

$= (o;\ m) :: \vec{x}.\ ((t_1/¿) \parallel \varepsilon)$

$= \{\ (o;\ m) :: \vec{x}.\ (t_1/¿)\ \}$

$= (o;\ m) :: \vec{x}.\ (t_1/¿) \parallel \varepsilon$

$= (((o;\ m) :: \vec{x}.\ t_1)\ /¿) \parallel (((¿;\ m) :: \vec{x}.\ t_2)\ /¿)\ .$

If it is given by a commuting case, then observe that $(t_2/¿)/(m \propto \vec{x}) = t_2/¿$ because $\mathrm{cc}(m)$ does not appear free in t_2. Using the definition of $\parallel$ in the commuting case, we then compute:

$(o;\ m) :: \vec{x}.\ ((t_1/¿) \parallel (t_2/¿))$

$= \{\ (o;\ m) :: \vec{x}.\ t \mid t \in ((t_1/¿) \parallel (t_2/¿))\ \}$

$= \{\ (o;\ m) :: \vec{x}.\ t \mid (t_2/¿)/(m \propto \vec{x}) \simeq t_2' \wedge t \in t_1/¿ \parallel t_2'\ \}$

$\subseteq ((o;\ m) :: \vec{x}.\ (t_1/¿)) \parallel t_2/¿$

$= (((o;\ m) :: \vec{x}.\ t_1)\ /¿) \parallel (((¿;\ m) :: \vec{x}.\ t_2)\ /¿)\ .$

This case analysis is exhaustive, so we conclude the result. □

LEMMA C.6. *If t is safely constrained, then $(t/(m \propto \vec{x}))/_{\textit{¿}}$, $(t/_{\textit{¿}})/(m \propto \vec{x})$, and $t/_{\textit{¿}}$ are all defined and equal whenever $t/(m \propto \vec{x})$ is defined.*

PROOF OF LEMMA C.6. The hypothesis that t is safely constrained is essential. Indeed, consider the unsafely constrained trace $t = (\textit{¿};\ \text{close on } a) :: (?;\ \text{close on } a) :: \varepsilon$. Then $t/\text{close on } a$ is defined but $(t/_{\textit{¿}})/\text{close on } a = ((?;\ \text{close on } a) :: \varepsilon)/\text{close on } a$ is undefined.

Assume that t is safely constrained and that $t/(m \propto \vec{x})$ is defined. It is immediate that traces $(t/(m \propto \vec{x}))/_{\textit{¿}}$ and $t/_{\textit{¿}}$ are defined. We proceed by induction on the length of t to show that $(t/(m \propto \vec{x}))/_{\textit{¿}} = (t/_{\textit{¿}})/(m \propto \vec{x}) = t/_{\textit{¿}}$. The base case occurs when $t = \varepsilon$, in which case the result is immediate. Assume next that $t = (s';\ m') :: \vec{y} . t'$. We consider two subcases, based on whether or not $\text{cc}(m) = \text{cc}(m')$.

Assume first that $\text{cc}(m) = \text{cc}(m')$. In this case, the fact that $t/(m \propto \vec{x})$ is defined implies $s' = \textit{¿}$. Then, where the middle equality is given by the induction hypothesis:

$$(t/(m \propto \vec{x}))/_{\textit{¿}} = t'/_{\textit{¿}} = (t'/_{\textit{¿}})/(m \propto \vec{x}) = (t/_{\textit{¿}})/(m \propto \vec{x}),$$

and $(t/(m \propto \vec{x}))/_{\textit{¿}} = t'/_{\textit{¿}} = t/_{\textit{¿}}$. These transitively give the equalities we wished to show.

Assume otherwise that $\text{cc}(m) \neq \text{cc}(m')$. We consider two subcases: if $s' = \textit{¿}$ or $s' \neq \textit{¿}$. If $s' = \textit{¿}$, then:

$$(t/(m \propto \vec{x}))/_{\textit{¿}}$$

by definition:

$$= ((s';\ m') :: \vec{y} . (t'/(m \propto \vec{x})))/_{\textit{¿}}$$
$$= (t'/(m \propto \vec{x}))/_{\textit{¿}}$$

by the induction hypothesis:

$$= (t'/_{\textit{¿}})/(m \propto \vec{x})$$
$$= (t/_{\textit{¿}})/(m \propto \vec{x}),$$

which establishes the first equality. The second equality is similarly established by appealing to the induction hypothesis:

$$(t/(m \propto \vec{x}))/_{\textit{¿}} = (t'/(m \propto \vec{x}))/_{\textit{¿}} = t'/_{\textit{¿}} = t/_{\textit{¿}}.$$

If $s' \neq \textit{¿}$, then the result follows by a similarly straightforward analysis. □

PROPOSITION 5.5. *Given safely constrained traces t_1 and t_2, $(t_1 \parallel t_2)/_{\textit{¿}} \subseteq (t_1/_{\textit{¿}}) \parallel (t_2/_{\textit{¿}})$.*

PROOF OF PROPOSITION 5.5. By lexicographic induction on the lengths of traces t_1 and t_2. The base cases are clear. If $t_1 = \varepsilon$, then

$$(\varepsilon \parallel t_2)/_{\textit{¿}} = \{\, t_2 \,\}/_{\textit{¿}} = \{\, t_2/_{\textit{¿}} \,\} = \varepsilon \parallel (t_2/_{\textit{¿}}) = (\varepsilon/_{\textit{¿}}) \parallel (t_2/_{\textit{¿}}).$$

The analysis for $t_2 = \varepsilon$ is symmetric.

Assume next that both traces are non-empty. We proceed by case analysis on which clause in the definition determines $t_1 \parallel t_2$.

The first clause is the synchronizing case in which the two traces send and receive the same message:

$$t_1 \parallel t_2 = ((!;\ m) :: \vec{x} . t_1') \parallel ((?;\ m) :: \vec{x} . t_2') = (\sigma;\ m) :: \vec{x} . \left(t_1' \parallel t_2' \right)$$

By the induction hypothesis, $\left(t'_1 \parallel t'_2\right) /_{¿} \subseteq \left(t'_1/_{¿}\right) \parallel \left(t_2/_{¿}'\right)$. By definition of $\cdot/_{¿}$ and monotonicity:

$$(t_1 \parallel t_2)/_{¿}$$
$$= \left((\sigma;\, m) :: \vec{x}.\, \left(t'_1 \parallel t'_2\right)\right) /_{¿}$$
$$= (\sigma;\, m) :: \vec{x}.\, \left(\left(t'_1 \parallel t'_2\right)/_{¿}\right)$$
$$\subseteq (\sigma;\, m) :: \vec{x}.\, \left(\left(t'_1/_{¿}\right) \parallel \left(t'_2/_{¿}\right)\right)$$
$$= \left(\left((!;\, m) :: \vec{x}.\, t'_1\right)/_{¿}\right) \parallel \left(\left((?;\, m) :: \vec{x}.\, t'_2\right)/_{¿}\right)$$
$$= (t_1/_{¿}) \parallel (t_2/_{¿}).$$

This is what we wanted to show.

The second clause is the principal case in which, without loss of generality, the head of t_1 is (s, m_1) and the head of t_2 is $(¿, m_2)$, where both m_1 and m_2 have the same carrier channel. It has two subcases. In the first subcase, $m_1 \neq m_2$ or the two traces bind different numbers of variables, and $t_1 \parallel t_2$ is empty. It follows that $(t_1 \parallel t_2)/_{¿}$ is empty, so the result holds trivially. The second subcase is when the interleaving is defined:

$$\left((s;\, m) :: \vec{x}.\, t'_1\right) \parallel \left((¿;\, m) :: \vec{x}.\, t'_2\right) = (s;\, m) :: \vec{x}.\, \left(t'_1 \parallel t'_2\right).$$

Assume first that $s = ¿$, then:

$$\left(\left((s;\, m) :: \vec{x}.\, t'_1\right) \parallel \left((¿;\, m) :: \vec{x}.\, t'_2\right)\right) /_{¿}$$
$$= \left((s;\, m) :: \vec{x}.\, \left(t'_1 \parallel t'_2\right)\right) /_{¿}$$
$$= \left(t'_1 \parallel t'_2\right) /_{¿}$$

which by the induction hypothesis,

$$\subseteq t'_1/_{¿} \parallel t'_2/_{¿}$$
$$= (t_1/_{¿}) \parallel (t_2/_{¿}).$$

Assume otherwise that $s \neq ¿$: We then compute, using the induction hypothesis, monotonicity, and the definition of $\parallel$:

$$\left(\left((s;\, m) :: \vec{x}.\, t'_1\right) \parallel \left((¿;\, m) :: \vec{x}.\, t'_2\right)\right) /_{¿}$$
$$= \left((s;\, m) :: \vec{x}.\, \left(t'_1 \parallel t'_2\right)\right) /_{¿}$$
$$= (s;\, m) :: \vec{x}.\, \left(\left(t'_1 \parallel t'_2\right) /_{¿}\right)$$
$$\subseteq (s;\, m) :: \vec{x}.\, \left(\left(t'_1/_{¿}\right) \parallel \left(t'_2/_{¿}\right)\right)$$

which by lemma C.5:

$$\subseteq (t_1/_{¿}) \parallel (t_2/_{¿}).$$

This completes the analysis of the synchronizing cases.

Next, we show the result for the commuting cases, where the heads of each trace perform actions on different channels. In this case, $t_i = (s_i;\, m_i) :: \vec{x}_i.\, t'_i$ for $1 \leq i \leq 2$ with $\mathrm{cc}(m_1) \neq \mathrm{cc}(m_2)$. We compute:

$$(t_1 \parallel t_2)/_{¿}$$
$$= \bigcup_{1 \leq i \neq j \leq 2} \left\{\, (s_i;\, m_i) :: \vec{x}_i.\, t \mid t_j/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in t'_i \parallel T'_j \,\right\}/_{¿}$$

using the fact the operation $\cdot/\xi$ is idempotent:

$$= \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, (t/\xi) \mid t_j/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in t'_i \parallel T'_j \,\}/\xi$$

$$= \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid t_j/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in (t'_i \parallel T'_j)/\xi \,\}/\xi$$

by the induction hypothesis:

$$= \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid t_j/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in (t'_i/\xi) \parallel (T'_j/\xi) \,\}/\xi$$

$$= \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid (t_j/(m_i \propto \vec{x}_i))/\xi \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}/\xi$$

Let X be the last set in this sequence. We proceed by case analysis on s_1 and s_2.

If $s_1 = s_2 = \xi$, then by lemma C.6, X is:

$$X \subseteq \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid (t_j/\xi)/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}/\xi$$

$$\subseteq \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid t_j/\xi \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}/\xi$$

$$= \bigcup_{1 \le i \ne j \le 2} \{\, t \mid t_j/\xi \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}$$

$$= \bigcup_{1 \le i \ne j \le 2} \{\, t \mid \in (t'_i/\xi) \parallel (t'_j/\xi) \,\}$$

$$= (t'_1/\xi) \parallel (t'_2/\xi)$$

$$= (t_1/\xi) \parallel (t_2/\xi),$$

which is what we wanted to show.

If neither of s_1 or s_2 is ξ, then by lemma C.6, X is:

$$X \subseteq \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid (t_j/\xi)/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}/\xi$$

$$= \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid (t_j/\xi)/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}$$

$$= \bigcup_{1 \le i \ne j \le 2} \{\, (s_i;\, m_i) :: \vec{x}_i \,.\, t \mid ((s_j;\, m_j) :: \vec{x}_j \,.\, (t'_j/\xi))/(m_i \propto \vec{x}_i) \simeq T'_j \wedge t \in (t'_i/\xi) \parallel T'_j \,\}$$

$$= ((s_1;\, m_1) :: \vec{x}_1 \,.\, (t'_1/\xi)) \parallel ((s_2;\, m_2) :: \vec{x}_2 \,.\, (t'_2/\xi))$$

$$= (t_1/\xi) \parallel (t_2/\xi).$$

If, without loss of generality, $s_1 = \xi$ and $s_2 \ne \xi$, then X is:

$$X = \{\, (\xi;\, m_1) :: \vec{x}_1 \,.\, t \mid (t_2/(m_1 \propto \vec{x}_1))/\xi \simeq T'_2 \wedge t \in (t'_1/\xi) \parallel T'_2 \,\}/\xi \ \cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/(m_2 \propto \vec{x}_2))/\xi \simeq T'_1 \wedge t \in (t'_2/\xi) \parallel T'_1 \,\}/\xi,$$

which by lemma C.6:

$$\subseteq \{\, (\zeta;\, m_1) :: \vec{x}_1 \,.\, t \mid t_2/\zeta \simeq T_2' \wedge t \in (t_1'/\zeta) \parallel T_2' \,\}/\zeta \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/(m_2 \propto \vec{x}_2))/\zeta \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}/\zeta$$
$$= \{\, t \mid t_2/\zeta \simeq T_2' \wedge t \in (t_1'/\zeta) \parallel T_2' \,\} \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/(m_2 \propto \vec{x}_2))/\zeta \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}/\zeta$$
$$= \{\, t \mid t \in (t_1'/\zeta) \parallel (t_2/\zeta) \,\} \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/(m_2 \propto \vec{x}_2))/\zeta \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}/\zeta$$
$$= (t_1'/\zeta) \parallel (t_2/\zeta) \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/(m_2 \propto \vec{x}_2))/\zeta \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}/\zeta$$
$$= (t_1'/\zeta) \parallel (t_2/\zeta) \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/(m_2 \propto \vec{x}_2))/\zeta \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\},$$

which by lemma C.6:

$$= (t_1'/\zeta) \parallel (t_2/\zeta) \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1/\zeta)/(m_2 \propto \vec{x}_2) \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}$$
$$= (t_1'/\zeta) \parallel (t_2/\zeta) \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1'/\zeta)/(m_2 \propto \vec{x}_2) \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}$$

At this point, observe that if T_1' is defined, then the carrier of the head of t_1'/ζ, if it exists, cannot be $\mathrm{cc}(m_2)$. Indeed, the head of t_1'/ζ has sign ?, ! or σ, and if its carrier were $\mathrm{cc}(m_2)$, then the reduction by $(m_2 \propto \vec{x}_2)$ would be undefined. By lemma C.6, this last union is:

$$= (t_1'/\zeta) \parallel (t_2/\zeta) \,\cup$$
$$\{\, (s_2;\, m_2) :: \vec{x}_2 \,.\, t \mid (t_1'/\zeta)/(m_2 \propto \vec{x}_2) \simeq T_1' \wedge t \in (t_2'/\zeta) \parallel T_1' \,\}/\zeta$$

By the just-above observation, the following interleaving must be given by a commuting case:

$$\subseteq (t_1'/\zeta) \parallel (t_2/\zeta) \,\cup$$
$$t_1'/\zeta \parallel ((s_2;\, m_2) :: \vec{x}_2 \,.\, (t_2'/\zeta))$$
$$= (t_1'/\zeta) \parallel (t_2/\zeta)$$
$$= (t_1/\zeta) \parallel (t_2/\zeta).$$

This case analysis is exhaustive and completes the proof.   □

PROPOSITION 5.6. *If* $t \in [\![\, \Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I}, \Gamma \vdash a : A) \,]\!]$ *and* $\Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I} \vdash a : A)$ *is well-formed, then* $(t \setminus \mathrm{dom}(\Gamma)) \in [\![\, \Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I} \vdash a : A) \,]\!]$.

PROOF OF PROPOSITION 5.6. Recall that $t \in [\![\, \Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I}, \Gamma \vdash a : A) \,]\!]$ is inductively defined in section 4.1 using a subset relationship. We proceed by induction on this derivation.

The base case, induced by $[\![\, \Pi \# \cdot \,]\!] \supseteq \{\, \varepsilon \,\}$ holds vacuously.

For the inductive steps, we show several representative cases.

Assume that $t \in [\![\, \Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I}, \Gamma \vdash a : \oplus_{l \in L} A_l) \,]\!]$ because $t = (!;\ \text{label } k \text{ on } a) :: t'$ with $t' \in [\![\, \Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid (\mathrm{I}, \Gamma)/\pi \vdash a : A_k) \,]\!]$, where $\pi = \text{label } k$ on a. By the induction hypothesis, $(t' \setminus \mathrm{dom}(\Gamma)) \in [\![\, \Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid \mathrm{I}/\pi \vdash a : A_k) \,]\!]$. Then by the same clause, $(t \setminus \mathrm{dom}(\Gamma)) \in [\![\, \Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I} \vdash a : \oplus_{l \in L} A_l) \,]\!]$.

Assume that $t \in [\![\, \Pi \# \mathcal{G} \mid (\Delta \mid \mathrm{I}, \Gamma, a : \oplus_{l \in L} A_l \vdash c : C) \,]\!]$ because $t = (\sigma;\ \text{label } k \text{ on } a) :: t'$ with $t' \in [\![\, \Pi/\pi \# \mathcal{G}/\pi \mid (\Delta/\pi \mid (\mathrm{I}, \Gamma)/\pi, a : A_k \vdash c : C/\pi) \,]\!]$, where $\pi = \text{label } k$ on a. This case involves

two subcases: the first where a is deleted with Γ, and where a is not deleted. The second subcase is analogous to the previous case, so assume that a is deleted. We must show that

$$t \setminus \text{dom}(\Gamma, a : \oplus_{l \in L} A_l) \in [\![\Pi \;\#\; \mathcal{G} \mid (\Delta \mid \text{I} \vdash c : C)]\!].$$

Observe that the specifications

$$\Pi/\pi \;\#\; \mathcal{G}/\pi \mid (\Delta/\pi \mid (\text{I}, \Gamma)/\pi, a : A_k \vdash c : C/\pi) \quad \text{and} \quad \Pi \;\#\; \mathcal{G} \mid (\Delta \mid \text{I}, \Gamma, a : A_k \vdash c : C)$$

are syntactically equal by the well-formedness assumption (reduction by an observed communication behaves as the identity on types that do not mention the communication's carrier channel). By the induction hypothesis,

$$t' \setminus \text{dom}(\Gamma, a : A_k) \in [\![\Pi/\pi \;\#\; \mathcal{G}/\pi \mid (\Delta/\pi \mid \text{I}/\pi \vdash c : C/\pi)]\!]$$

By the above syntactic observation,

$$t' \setminus \text{dom}(\Gamma, a : A_k) \in [\![\Pi \;\#\; \mathcal{G} \mid (\Delta \mid \text{I} \vdash c : C)]\!].$$

By definition of the deletion operation,

$$t \setminus \text{dom}(\Gamma, a : \oplus_{l \in L} A_l) \in [\![\Pi \;\#\; \mathcal{G} \mid (\Delta \mid \text{I} \vdash c : C)]\!].$$

The remaining cases are analogous.   □

PROPOSITION 5.7. *If* $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \;\#\; \mathcal{G}_1]\!]$ *and* $t_2 \in [\![\Pi, \overline{\mathcal{G}_1} \;\#\; \mathcal{G}_2]\!]$, *then* $t_1 \parallel t_2 \subseteq [\![\Pi \;\#\; \mathcal{G}_1 \mid \mathcal{G}_2]\!]$.

PROOF OF PROPOSITION 5.7. By lexicographic induction on the lengths of t_1 and t_2. The base cases where either trace is ε are obvious. Assume next that $t_i = (s_i;\ m_i) :: \vec{x}_i \,.\, \tau_i$ for $i = 1, 2$, and set $\pi_i = (m_i \propto \vec{x}_i)$. We proceed by case analysis on which clause defines $t_1 \parallel t_2$.

The first synchronizing case, where $m_1 = m_2$ and $\vec{x}_1 = \vec{x}_2$ and, without loss of generality, $s_1 = \,!$ and $s_2 = ?$, is impossible. Indeed, it would require $\mathcal{G}_1$ and $\mathcal{G}_2$ to type the same channel, but this would render the specification $\Pi \;\#\; \mathcal{G}_1 \mid \mathcal{G}_2$ ill-formed.

The second synchronizing case, where $m_1 = m_2$ and $\vec{x}_1 = \vec{x}_2$, and without loss of generality, $s_2 = \mathord{\text{¿}}$ is possible. We proceed by case analysis on s_1 and m_1. If $s_1 = ?$ and $m_1 = \text{label } a \text{ on } k$, then

(1) $t_1 = (!;\ \text{label } a \text{ on } k) :: \tau_1$
(2) $t_2 = (\text{¿};\ \text{label } a \text{ on } k) :: \tau_2$

Inversion on the derivations $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \;\#\; \mathcal{G}_1]\!]$ and $t_2 \in [\![\Pi, \overline{\mathcal{G}_1} \;\#\; \mathcal{G}_2]\!]$ gives:

(3) $\mathcal{G}_1 = \mathcal{G}_1' \mid (\Delta_1 \mid \text{I}_1 \vdash a : \oplus_{l \in L} A_l)$
(4) $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \;\#\; \mathcal{G}_1' \mid (\Delta_1 \mid \text{I}_1 \vdash a : \oplus_{l \in L} A_l)]\!]$
(5) $t_2 \in [\![\Pi, \overline{\mathcal{G}_1', \Delta_1, \text{I}_1}, a : \oplus_{l \in L} A_l \;\#\; \mathcal{G}_2]\!]$
(6) $\tau_1 \in [\![(\Pi, \overline{\mathcal{G}_2})/\pi_1 \;\#\; \mathcal{G}_1'/\pi_1 \mid (\Delta_1/\pi_1 \mid \text{I}_1/\pi_1 \vdash a : A_k)]\!]$
(7) $\tau_2 \in [\![(\Pi, \overline{\mathcal{G}_1', \Delta_1, \text{I}_1})/\pi_1, a : A_k \;\#\; \mathcal{G}_2/\pi_1]\!]$

By the induction hypothesis:

(8) $\tau_1 \parallel \tau_2 \subseteq [\![\Pi/\pi_1 \;\#\; \mathcal{G}_1'/\pi_1 \mid (\Delta_1/\pi_1 \mid \text{I}_1/\pi_1 \vdash a : A_k) \mid \mathcal{G}_2/\pi_1]\!]$

Then by definition of trace sets for specifications,

(9) $(!;\ \text{label } a \text{ on } k) :: (\tau_1 \parallel \tau_2) \subseteq [\![\Pi \;\#\; \mathcal{G}_1' \mid (\Delta_1 \mid \text{I}_1 \vdash a : \oplus_{l \in L} A_l) \mid \mathcal{G}_2]\!]$

But this is exactly the desired inclusion $t_1 \parallel t_2 \subseteq [\![\Pi \;\#\; \mathcal{G}_1 \mid \mathcal{G}_2]\!]$. The analysis is analogous for the other subcases $s_1 \in \{\,!, \sigma, \text{¿}\,\}$. It is also analogous for other values of m_1.

Finally, we consider the commuting cases where $cc(m_1) \neq cc(m_2)$. In this case,

$$(s_1;\ m_1) :: \vec{x}_1 . t_1 \parallel (s_2;\ m_2) :: \vec{x}_2 . t_2$$

$$= \{\ (s_1;\ m_1) :: \vec{x}_1 . t \mid t_2/(m_1 \propto \vec{x}_1) \simeq T_2' \wedge t \in t_1 \parallel T_2'\ \} \cup$$

$$\cup\ \{\ (s_2;\ m_2) :: \vec{x}_2 . t \mid t_1/(m_2 \propto \vec{x}_2) \simeq T_1' \wedge t \in T_1' \parallel t_2\ \}$$

We must show that each of these two subsets is contained in $[\![\Pi \ \# \ \mathcal{G}_1 \mid \mathcal{G}_2]\!]$. We do so for the first; the second inclusion will follow by symmetry. If T_2' is undefined, then the inclusion is immediate (the set is empty). If T_2' is defined, then we consider two subcases: whether $s_1 = \text{¿}$ or $s_1 \neq \text{¿}$.

If $s_1 = \text{¿}$, then by inversion on $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \ \# \ \mathcal{G}_1]\!]$:

(1) $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \ \# \ \mathcal{G}_1]\!]$,
(2) $t_1' \in [\![\Pi', \overline{\mathcal{G}_2'} \ \# \ \mathcal{G}_1/\pi_1]\!]$ for some $\Pi', \overline{\mathcal{G}_2'}$ such that
(3) $\Pi, \overline{\mathcal{G}_2} \overset{\pi_1}{\Longrightarrow} \Pi', \overline{\mathcal{G}_2'}$.

That $T_2' = t_2/\pi_1$ is defined implies that $cc(\pi_1) \notin dom(\mathcal{G}_2)$, so $cc(\pi_1) \in dom(\Pi)$. It follows that $\Pi, \overline{\mathcal{G}_2} \overset{\pi_1}{\Longrightarrow} \Pi', \overline{\mathcal{G}_2'}$ because:

(4) $\Pi \overset{\pi_1}{\Longrightarrow} \Pi'$ and

(5) $\overline{\mathcal{G}_2'} = \mathcal{G}_2/\pi_1$.

These facts jointly imply

(6) $\Pi, \overline{\mathcal{G}_2} \ \# \ \mathcal{G}_1 \overset{(\text{¿},\pi_1)}{\Longrightarrow} \Pi', \overline{\mathcal{G}_2/\pi_1} \ \# \ \mathcal{G}_1/\pi_1$.

By proposition C.4, $T_2' \in [\![\Pi', \overline{\mathcal{G}_2/\pi_1} \ \# \ \mathcal{G}_1/\pi_1]\!]$. By the induction hypothesis,

$$t_1' \parallel T_2' \subseteq [\![\Pi' \ \# \ \mathcal{G}_1/\pi_1 \mid \mathcal{G}_2/\pi_1]\!].$$

By the same clause giving $t_1 \in [\![\Pi, \overline{\mathcal{G}_2} \ \# \ \mathcal{G}_1]\!]$, we deduce

$$(\text{¿};\ m_1) :: \vec{x}_1 . (t_1' \parallel T_2') \subseteq [\![\Pi' \ \# \ \mathcal{G}_1/\pi_1 \mid \mathcal{G}_2/\pi_1]\!],$$

which is what we wanted to show.

The case where $s_1 \neq \text{¿}$ is analogous. This completes the case analyses so we conclude the result.   □

PROPOSITION 5.8. *If* $t_1 \in [\![\Pi, \Delta_2, I_2, c : C \ \# \ (\Delta_1 \mid I_1 \vdash a : A)]\!]$ *and* $t_2 \in [\![\Pi, \Delta_1, I_1 \ \# \ (\Delta_2, a : A \mid I_2 \vdash c : C)]\!]$, *then* $t_1 \parallel t_2 \subseteq [\![\Pi \ \# \ (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C)]\!]$.

PROOF OF PROPOSITION 5.8. By induction on definition of $t_1 \parallel t_2$.

The base cases where traces t_1 or t_2 are ε are impossible by inversion on the derivations of $t_1 \in [\![\Pi, \Delta_2, I_2, c : C \ \# \ (\Delta_1 \mid I_1 \vdash a : A)]\!]$ and $t_2 \in [\![\Pi, \Delta_1, I_1 \ \# \ (\Delta_2, a : A \mid I_2 \vdash c : C)]\!]$.

In the inductive cases, $t_i = (s_i;\ m_i) :: \vec{x}_i . t_i'$.

Assume first that we are in the principal synchronizing case: $cc(m_1) = cc(m_2)$ and $\vec{x}_1 = \vec{x}_2$, and without loss of generality, $s_1 = !$ and $s_2 = ?$. We proceed by case analysis on m_1. If $m_1 = $ label a on k, then

(1) $t_1 = (!;\ \text{label } a \text{ on } k) :: \tau_1$
(2) $t_2 = (?;\ \text{label } a \text{ on } k) :: \tau_2$
(3) $A = \oplus_{l \in L} A_l$ with $k \in L$.

Set $\pi = (\text{label } a \text{ on } k \propto \emptyset)$. Inversion on the derivations $t_1 \in [\![\Pi, \Delta_2, I_2, c : C \ \# \ (\Delta_1 \mid I_1 \vdash a : A)]\!]$ and $t_2 \in [\![\Pi, \Delta_1, I_1 \ \# \ (\Delta_2, a : A \mid I_2 \vdash c : C)]\!]$ gives:

(4) $\tau_1 \in [\![(\Pi, \Delta_2, I_2, c : C)/\pi \ \# \ (\Delta_1/\pi \mid I_1/\pi \vdash a : A_k)]\!]$

(5) $\tau_2 \in [\![(\Pi, \Delta_1, I_1)/\pi \; \# \; (\Delta_2/\pi, a : A_k \mid I_2/\pi \vdash c : C/\pi)]\!]$

By the induction hypothesis, $\tau_1 \parallel \tau_2 \subseteq [\![\Pi/\pi \; \# \; (\Delta_1, \Delta_2)/\pi \mid (I_1, I_2)/\pi, a : A_k \vdash c : C/\pi]\!]$. It then follows that $(\sigma; \; \text{label } a \text{ on } k) :: (\tau_1 \parallel \tau_2) \subseteq [\![\Pi \; \# \; (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C)]\!]$. But this is exactly what we wanted to show, for $t_1 \parallel t_2 = (\sigma; \; \text{label } a \text{ on } k) :: (\tau_1 \parallel \tau_2)$. The case analyses for other choices of m_1 are analogous.

Assume next that it is the synchronizing case dealing with constraints. In particular, $m_1 = m_2$, $\vec{x}_1 = \vec{x}_2$, and assume that $s_2 = \textrm{¿}$ (this is not without loss of generality, but the case where $s_1 = \textrm{¿}$ is analogous). Set $\pi = (m_1 \propto \vec{x}_1)$. By inversion on the elementhood relation for t_2, we know that $\text{cc}(m_2) \neq a$. We proceed by case analysis on s_1 and m_1. Assume first that $s_1 = ?$ and $m_1 = \text{label } k \text{ on } b$. Then by inversion on elementhood for t_1, we have $\Delta_1 = \Delta, b : \oplus_{l \in L} B_l$ with $k \in L$, and by inversion on elementhood for both traces:

(1) $t_1 \in [\![\Pi, \Delta_2, I_2, c : C \; \# \; (\Delta_1 \mid I_1 \vdash a : A)]\!]$
(2) $t_1' \in [\![(\Pi, \Delta_2, I_2, c : C)/\pi \; \# \; (\Delta/\pi, b : B_k \mid I_1/\pi \vdash a : A/\pi)]\!]$
(3) $t_2 \in [\![\Pi, \Delta_1, I_1 \; \# \; (\Delta_2, a : A \mid I_2 \vdash c : C)]\!]$
(4) $t_2' \in [\![(\Pi, I_1, \Delta)/\pi, b : B_k \; \# \; (\Delta_2, a : A \mid I_2 \vdash c : C)/\pi]\!]$

By the induction hypothesis,

$$ t_1' \parallel t_2' \subseteq [\![\Pi/\pi \; \# \; ((\Delta, \Delta_2)/\pi, b : B_k \mid (I_1, I_2, a : A)/\pi \vdash c : C/\pi)]\!]. $$

By the same clause giving elementhood for t_1,

$$ t_1 \parallel t_2 = (s_1; \; m_1) :: \vec{x}_1 . (t_1' \parallel t_2') \subseteq [\![\Pi \; \# \; (\Delta_1, \Delta, b : \oplus_{l \in L} B_l \mid I_1, a : A, I_2 \vdash c : C)]\!]. $$

This is exactly what we wanted to show. The case analyses for other choices of s_1 and m_1 are equally mundane.

Finally, assume that $t_1 \parallel t_2$ is given by a commuting case, where $\text{cc}(m_1) \neq \text{cc}(m_2)$. In this case,

$$ (s_1; \; m_1) :: \vec{x}_1 . t_1' \parallel (s_2; \; m_2) :: \vec{x}_2 . t_2' $$
$$ = \{ \, (s_1; \; m_1) :: \vec{x}_1 . t \mid t_2/(m_1 \propto \vec{x}_1) \simeq T_2' \wedge t \in t_1' \parallel T_2' \, \} \; \cup $$
$$ \cup \, \{ \, (s_2; \; m_2) :: \vec{x}_2 . t \mid t_1/(m_2 \propto \vec{x}_2) \simeq T_1' \wedge t \in T_1' \parallel t_2' \, \} $$

We must show that each of these two subsets is contained in $[\![\Pi \; \# \; (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C)]\!]$. We do so for the first subset; the second inclusion is analogous. If T_2' is undefined, then the inclusion is immediate (the set is empty). Assume next that T_2' is defined; this implies that $\text{cc}(m_1) \neq a$. Again, we proceed by case analysis on s_1 and m_1, and give an illustrative case. Assume first that $s_1 = \sigma$ and $m_1 = \text{label } k \text{ on } b$. By inversion on elementhood for t_1,

(1) $I_1 = I_1', b : \odot_{l \in L} B_l$ with $k \in L$ and $\odot \in \{ \oplus, \& \}$;
(2) $t_1' \in [\![(\Pi, \Delta_2, I_2, c : C)/\pi \; \# \; (\Delta_1/\pi \mid I_1'/\pi, b : B_k \vdash a : A/\pi)]\!]$.

By proposition C.4,

(3) $T_2' \in [\![(\Pi, \Delta_1, I_1')/\pi, b : B_k \; \# \; (\Delta_2, a : A \mid I_2 \vdash c : C)/\pi]\!]$.

By the induction hypothesis,

$$ t_1' \parallel T_2' \subseteq [\![\Pi/\pi \; \# \; ((\Delta_1, \Delta_2)/\pi \mid (I_1', I_2, a : A)/\pi, b : B_k \vdash c : C/\pi)]\!]. $$

By the same clause giving the elementhood relationship for t_1, we deduce

$$ (s_1; \; m_1) :: \vec{x}_1 . (t_1' \parallel T_2') \subseteq [\![\Pi \; \# \; (\Delta_1, \Delta_2 \mid I_1', b : \odot_{l \in L} B_l, I_2, a : A \vdash c : C)]\!], $$

which is what we wanted to show. The remaining cases are analogous. This completes the case analyses so we conclude the result. □

THEOREM 7.1 (SOUNDNESS). *The typechecking algorithm is sound: if any of*

(1) $P \Vdash \Pi \, \# \, (\Delta \mid I \vdash a : A) \mathbin{/\!\!/} \mathfrak{T}$,

(2) $P \Rightarrow \Pi \, \# \, (\Delta \mid I \vdash [a : A]) \mathbin{/\!\!/} \mathfrak{T}$, or

(3) $P \Rightarrow \Pi \, \# \, (\Delta, [a : A] \mid I \vdash c : C) \mathbin{/\!\!/} \mathfrak{T}$,

then $\mathfrak{T}/_{\vec{\iota}} = [\![P]\!]$ (the set $\mathfrak{T}$ extents the executions of P with constraints) and $\mathfrak{T} \subseteq [\![\Pi \, \# \, \Delta \mid I \vdash a : A]\!]$ (the constraint set $\mathfrak{T}$ satisfies the specification $\Pi \, \# \, \Delta \mid I \vdash a : A$).

PROOF OF THEOREM 7.1. By induction on the derivation.

CASE (NEW). The result is immediate by the induction hypothesis, proposition 5.6, and eq. (8).

CASE (PAR). Recall

$$\frac{P \Vdash \Pi, \Delta_2, I_2, c : C \, \# \, (\Delta_1 \mid I_1 \vdash a : A) \mathbin{/\!\!/} \mathfrak{T}_1 \quad Q \Vdash \Pi, \Delta_1, I_1 \, \# \, (\Delta_2, a : A \mid I_2 \vdash c : C) \mathbin{/\!\!/} \mathfrak{T}_2 \quad \mathfrak{T}_1 \therefore \mathfrak{T}_2}{P \parallel_a Q \Vdash \Pi \, \# \, (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C) \mathbin{/\!\!/} \mathfrak{T}_1 \parallel \mathfrak{T}_2} \text{ (PAR)}$$

By the induction hypothesis, $\mathfrak{T}_1/_{\vec{\iota}} = [\![P]\!]$ and $\mathfrak{T}_2/_{\vec{\iota}} = [\![Q]\!]$. By definition of $\mathfrak{T}_1 \therefore \mathfrak{T}_2$,

$$[\![P \parallel_a Q]\!] = [\![P]\!] \parallel [\![Q]\!] = (\mathfrak{T}_1/_{\vec{\iota}}) \parallel (\mathfrak{T}_2/_{\vec{\iota}}) = (\mathfrak{T}_1 \parallel \mathfrak{T}_2)/_{\vec{\iota}}.$$

The inclusion $\mathfrak{T}_1/_{\vec{\iota}} \parallel \mathfrak{T}_2/_{\vec{\iota}} \subseteq [\![\Pi \, \# \, (\Delta_1, \Delta_2 \mid I_1, a : A, I_2 \vdash c : C)]\!]$ is given by proposition 5.8.

CASE ($\Vdash$R). The result is immediate by the induction hypothesis.

CASE ($\Vdash$L). The result is immediate by the induction hypothesis.

CASE ($\Rightarrow$**1**R). Recall

$$\frac{}{\text{close } a \Rightarrow \Pi \, \# \, (\cdot \mid \cdot \vdash [a : \mathbf{1}]) \mathbin{/\!\!/} \{ \, (!; \text{ close on } a) :: \varepsilon \, \}} \text{ ($\Rightarrow$\textbf{1}R)}$$

The inclusion $\mathfrak{T} = \{ \, (!; \text{ close on } a) :: \varepsilon \, \} \subseteq [\![\Pi \, \# \, \cdot \mid \cdot \vdash a : \mathbf{1}]\!]$ is easily checked. Moreover, $\mathfrak{T} \setminus \text{fc}(\Pi) = \mathfrak{T} = [\![\text{close } a]\!]$ by eq. (1), so we are done.

CASE ($\Rightarrow\oplus$R). The equality is by the induction hypothesis and eq. (3). The inclusion is by the induction hypothesis and eq. (9).

CASE ($\Rightarrow$&R). The equality is by the induction hypothesis and eq. (4). The inclusion is by the induction hypothesis and eq. (22).

CASE ($\Rightarrow\oplus$L). Analogous to ($\Rightarrow$&R)

CASE ($\Rightarrow$&L). Analogous to ($\Rightarrow\oplus$R).

CASE ($\Rightarrow\otimes$R). By the induction hypothesis, $\mathfrak{T}_1/_{\vec{\iota}} = [\![P]\!]$ and $\mathfrak{T}_2/_{\vec{\iota}} = [\![Q]\!]$. By eq. (5),

$$[\![\text{send } a \, (b.P); \, Q]\!]$$
$$= \{ \, (!; \text{ chan on } a) :: b \, . \, t \mid t_P \in [\![P]\!], t_Q \in [\![Q]\!], t \in t_P \parallel t_Q \, \}$$
$$= \{ \, (!; \text{ chan on } a) :: b \, . \, t \mid t_P \in \mathfrak{T}_1/_{\vec{\iota}}, t_Q \in \mathfrak{T}_2/_{\vec{\iota}}, t \in t_P \parallel t_Q \, \}$$
$$= \{ \, (!; \text{ chan on } a) :: b \, . \, t \mid t \in \mathfrak{T}_1/_{\vec{\iota}} \parallel \mathfrak{T}_2/_{\vec{\iota}} \, \}$$

which by $\mathfrak{T}_1 \therefore \mathfrak{T}_2$,

$$= \{ \, (!; \text{ chan on } a) :: b \, . \, t \mid t \in (\mathfrak{T}_1 \parallel \mathfrak{T}_2)/_{\vec{\iota}} \, \}$$
$$= (!; \text{ chan on } a) :: b \, . \, ((\mathfrak{T}_1 \parallel \mathfrak{T}_2)/_{\vec{\iota}})$$
$$= ((!; \text{ chan on } a) :: b \, . \, (\mathfrak{T}_1 \parallel \mathfrak{T}_2))/_{\vec{\iota}}$$

The desired inclusion is given by the induction hypothesis, proposition 5.7, and eq. (16).

CASE ($\Rightarrow\multimap$R). Recall

$$\frac{P \Vdash \Pi/\pi \, \# \, (\Delta/\pi, b : B \mid I/\pi \vdash a : A) \mathbin{/\!\!/} \mathfrak{T} \quad \pi = \text{chan } b \text{ on } a}{b \leftarrow \text{receive } a; \, P \Rightarrow \Pi \, \# \, (\Delta \mid I \vdash [a : (b : B) \multimap (a : A)]) \mathbin{/\!\!/} (?; \text{ chan on } a) :: b \, . \, \mathfrak{T}} \text{ ($\Rightarrow\multimap$R)}$$

The equality is by the induction hypothesis and eq. (6). The inclusion is by the induction hypothesis and eq. (26).

CASE ($\Rightarrow\otimes$L). Analogous to ($\Rightarrow\multimap$R).

Case ($\Rightarrow\multimap$L). Analogous to ($\Rightarrow\otimes$R).

Case ($\Rightarrow\mathbf{1}$L). The desired inclusion is given by the induction hypothesis and eq. (13). The desired equality is given by the induction hypothesis and eq. (2).

Case ($\Rightarrow\mathbf{1}$[R]), ($\Rightarrow\mathbf{1}$[L]), ($\Rightarrow\oplus\&$[R]), ($\Rightarrow\oplus\&$[L]), ($\Rightarrow\otimes\multimap$[R]), or ($\Rightarrow\otimes\multimap$[L]). The inclusion is given by the induction hypothesis and the appropriate choice of eqs. (18) to (20). The equality is immediate by the induction hypothesis and the definition of $(-)/_{¿}$. $\qquad\square$

Lemma C.7. *If $t \in [\![\cdot \# \mathcal{G}]\!]$, then $t/_{¿} = t$.*

Proof of Lemma C.7. By induction on the derivation $t \in [\![\cdot \# \mathcal{G}]\!]$. $\qquad\square$

Corollary 7.2. *If $P \Vdash \cdot \# (\Delta \mid \mathrm{I} \vdash a : A) /\!\!/ \mathfrak{T}$, then $\mathfrak{T} = [\![P]\!]$ and $[\![P]\!] \subseteq [\![\cdot \# (\Delta \mid \mathrm{I} \vdash a : A)]\!]$.*

Proof of Corollary 7.2. Immediate by lemma C.7 and theorem 7.1. $\qquad\square$

Proposition 7.3 (Compositionality). *If $P \Vdash \cdot \# (\Delta_1 \mid \mathrm{I}_1 \vdash a : A) /\!\!/ \mathfrak{T}_1$ and $Q \Vdash \cdot \# (\Delta_2, a : A \mid \mathrm{I}_2 \vdash c : C) /\!\!/ \mathfrak{T}_2$, then $P \parallel_a Q \Vdash \cdot \# (\Delta_1, \Delta_2 \mid \mathrm{I}_1, a : A, \mathrm{I}_2 \vdash c : C) /\!\!/ \mathfrak{T}_1 \parallel \mathfrak{T}_2$.*

Proof of Proposition 7.3. Assume

- $P \Vdash \cdot \# (\Delta_1 \mid \mathrm{I}_1 \vdash a : A) /\!\!/ \mathfrak{T}_1$ and
- $Q \Vdash \cdot \# (\Delta_2, a : A \mid \mathrm{I}_2 \vdash c : C) /\!\!/ \mathfrak{T}_2$.

By induction on the two derivations, we can weaken them to

- $P \Vdash \Delta_2, \mathrm{I}_2, c : C \# (\Delta_1 \mid \mathrm{I}_1 \vdash a : A) /\!\!/ \mathfrak{T}_1$ and
- $Q \Vdash \Delta_1, \mathrm{I}_1 \# (\Delta_2, a : A \mid \mathrm{I}_2 \vdash c : C) /\!\!/ \mathfrak{T}_2$.

To show $P \parallel_a Q \Vdash \cdot \# (\Delta_1, \Delta_2 \mid \mathrm{I}_1, a : A, \mathrm{I}_2 \vdash c : C) /\!\!/ \mathfrak{T}_1 \parallel \mathfrak{T}_2$, it is sufficient by corollary 7.2 and (Par) to show that $\mathfrak{T}_1 \therefore \mathfrak{T}_2$, i.e., that

$$\forall t_1 \in \mathfrak{T}_1 . \forall t_2 \in \mathfrak{T}_2 . (t_1/_{¿} \parallel t_2/_{¿}) = (t_1 \parallel t_2) /_{¿}.$$

Indeed, if $\mathfrak{T}_1 \therefore \mathfrak{T}_2$, then $P \parallel_a Q \Vdash \cdot \# (\Delta_1, \Delta_2 \mid \mathrm{I}_1, a : A, \mathrm{I}_2 \vdash c : C) /\!\!/ \mathfrak{T}_1 \parallel \mathfrak{T}_2$ by (Par). We show that $\mathfrak{T}_1 \therefore \mathfrak{T}_2$. By corollary 7.2 and lemma C.7, $t_i/_{¿} = t_i$ for all $t_i \in \mathfrak{T}_i$ with $i = 1, 2$. We deduce

$$(t_1/_{¿} \parallel t_2/_{¿}) = t_1 \parallel t_2 = (t_1 \parallel t_2) /_{¿},$$

where the last equality stems from the fact that $\cdot/_{¿}$ is the identity on traces with no constraints, and interleaving two traces without constraints does not introduce any constraints. This is what we wanted to show. $\qquad\square$

REFERENCES

Robert Atkey. 2017. "Observed Communication Semantics for Classical Processes". In: *Programming Languages and Systems* (Lecture Notes in Computer Science) 10201. 26th European Symposium on Programming. ESOP 2017 (Uppsala, Sweden, Apr. 22–29, 2017). Ed. by Hongseok Yang. Springer-Verlag GmbH Germany, Berlin, Germany, 56–82. isbn: 978-3-662-54434-1. doi: 10.1007/978-3-662-54434-1_3.

Stephanie Balzer and Frank Pfenning. Sept. 2017. "Manifest Sharing With Session Types". *Proceedings of the ACM on Programming Languages*, 1, ICFP, (Sept. 2017), 37. doi: 10.1145/3110281.

Luís Caires and Frank Pfenning. 2010. "Session Types as Intuitionistic Linear Propositions". In: *CONCUR 2010 — Concurrency Theory* (Lecture Notes in Computer Science) 6269. 21st International Conference, CONCUR 2010 (Paris, France, Aug. 31–Sept. 3, 2010). Ed. by Paul Gastin and François Laroussinie. Springer-Verlag Berlin Heidelberg, 222–236. isbn: 978-3-642-15374-7. doi: 10.1007/978-3-642-15375-4_16.

Luís Caires, Frank Pfenning, and Bernardo Toninho. Mar. 2016. "Linear Logic Propositions As Session Types". *Mathematical Structures in Computer Science*, 26, 3, (Mar. 2016), 367–423. *Behavioural Types Part 2*. doi: 10.1017/s0960129514000218.

Iliano Cervesato and Andre Scedrov. Oct. 2009. "Relating State-Based and Process-Based Concurrency Through Linear Logic (Full-Version)". *Information and Computation*, 207, 10, (Oct. 2009), 1044–1077. *Special Issue: 13th Workshop on Logic, Language, Information and Computation (WoLLIC 2006)*. doi: 10.1016/j.ic.2008.11.006.

Silvia Crafa, Daniele Varacca, and Nobuko Yoshida. 2007. "Compositional Event Structure Semantics for the Internal π-Calculus". In: *CONCUR 2007 — Concurrency Theory* (Lecture Notes in Computer Science) 4703. 18th International Conference on Concurrency Theory. CONCUR 2007 (Lisbon, Portugal, Sept. 3–8, 2007). Ed. by Luís Caires and Vasco T. Vasconcelos. Springer-Verlag Berlin Heidelberg, 317–332. ISBN: 978-3-540-74407-8. DOI: 10.1007/978-3-540-74407-8_22.

Ankush Das and Frank Pfenning. Aug. 26, 2020. "Session Types With Arithmetic Refinements". In: *31st International Conference on Concurrency Theory (CONCUR 2020)* (Leibniz International Proceedings in Informatics (LIPIcs)). 31st International Conference on Concurrency Theory. CONCUR 2020 (Vienna, Austria (Virtual Conference), Sept. 1–4, 2020). Ed. by Igor Konnov and Laura Kovács. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Saarbrücken/Wadern, Germany, (Aug. 26, 2020). ISBN: 978-3-95977-160-3. DOI: 10.4230/LIPIcs.CONCUR.2020.13.

Pierre-Malo Deniélou and Nobuko Yoshida. 2011. "Dynamic Multirole Session Types". In: *POPL'11*. 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL'11 (Austin, Texas, Jan. 26–28, 2011). Association for Computing Machinery, Inc., New York, New York, 435–446. ISBN: 978-1-4503-0490-0. DOI: 10.1145/1926385.1926435.

Farzaneh Derakhshan and Frank Pfenning. May 10, 2022. "Circular Proofs As Session-Typed Processes: A Local Validity Condition". *Logical Methods in Computer Science*, 18, 2, (May 10, 2022), 8. DOI: 10.46298/lmcs-18(2:8)2022.

Thomas Ehrhard. Aug. 2018. "An Introduction To Differential Linear Logic: Proof-Nets, Models and Antiderivatives". *Mathematical Structures in Computer Science*, 28, 7, (Aug. 2018), 995–1060. *Differential Linear Logic, Nets and Other Quantitative and Parallel Approaches to Proof-Theory*. DOI: 10.1017/s0960129516000372.

Jérôme Fortier and Luigi Santocanale. Sept. 2, 2013. "Cuts for Circular Proofs: Semantics and Cut-Elimination". In: *Computer Science Logic 2013* (Leibniz International Proceedings in Informatics (LIPIcs)) 23. Computer Science Logic 2013. CSL'13 (Torino, Italy, Sept. 2–5, 2013). Ed. by Simona Ronchi Della Rocca. European Association for Computer Science Logic. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Saarbrücken/Wadern, Germany, (Sept. 2, 2013), 248–262. ISBN: 978-3-939897-60-6. DOI: 10.4230/LIPIcs.CSL.2013.248.

Murdoch Gabbay and Dan Ghica. Sept. 24, 2012. "Game Semantics in the Nominal Model". *Electronic Notes in Theoretical Computer Science*, 286, (Sept. 24, 2012), 173–189. *Proceedings of the 28th Conference on the Mathematical Foundations of Programming Semantics (MFPS XXVIII)*. DOI: 10.1016/j.entcs.2012.08.012.

Kohei Honda. 1993. "Types for Dyadic Interaction". In: *CONCUR'93* (Lecture Notes in Computer Science) 715. 4th International Conference on Concurrency Theory (Hildesheim, Germany, Aug. 23–26, 1993). Ed. by Eike Best. Springer-Verlag Berlin Heidelberg, Berlin, 509–523. ISBN: 978-3-540-47968-0. DOI: 10.1007/3-540-57208-2_35.

Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. "Language Primitives and Type Discipline for Structured Communication-Based Programming". In: *Programming Languages and Systems* (Lecture Notes in Computer Science) 1381. 7th European Symposium on Programming. ESOP'98 (Lisbon, Portugal, Mar. 28–Apr. 4, 1998). Ed. by Chris Hankin. Joint European Conferences on Theory and Practice of Software, ETAPS'98. Springer-Verlag Berlin Heidelberg, 122–138. ISBN: 978-3-540-69722-0. DOI: 10.1007/BFb0053567.

Kohei Honda, Nobuko Yoshida, and Marco Carbone. Mar. 3, 2016. "Multiparty Asynchronous Session Types". *Journal of the ACM*, 63, 1, (Mar. 3, 2016), 9. DOI: 10.1145/2827695.

Ryan Kavanagh. May 2022. "Fairness and Communication-Based Semantics for Session-Typed Languages". *Information and Computation*, 285, B, (May 2022), 104892. DOI: 10.1016/j.ic.2022.104892.

Ryan Kavanagh and Brigitte Pientka. Apr. 2024. "Message-Observing Sessions". *Proceedings of the ACM on Programming Languages*, 8, OOPSLA1, (Apr. 2024), 142. DOI: 10.1145/3649859.

Benjamin Pierce. 2002. *Types and Programming Languages*. The MIT Press, Cambridge, Massachusetts. xxi+623 pp. ISBN: 0-262-16209-1.

Pedro Rocha and Luís Caires. Aug. 2021. "Propositions-As-Types and Shared State". *Proceedings of the ACM on Programming Languages*, 5, ICFP, (Aug. 2021), 79. DOI: 10.1145/3473584.

Pedro Rocha and Luís Caires. 2023. "Safe Session-Based Concurrency With Shared Linear State". In: *Programming Languages and Systems* (Lecture Notes in Computer Science) 13990. 32nd European Symposium on Programming. ESOP 2023 (Paris, France, Apr. 22–27, 2023). Ed. by Thomas Wies. Springer, Cham, Switzerland, 421–450. ISBN: 978-3-031-30044-8. DOI: 10.1007/978-3-031-30044-8_16.

Siva Somayyajula and Frank Pfenning. 2022. "Type-Based Termination for Futures". In: *7th International Conference on Formal Structures for Computation and Deduction* (Leibniz International Proceedings in Informatics (LIPIcs)). 7th International Conference on Formal Structures for Computation and Deduction. FSCD 2022 (Haifa, Israel, Aug. 2–5, 2022). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Saarbrücken/Wadern, Germany. ISBN: 978-3-95977-233-4. DOI: 10.4230/LIPIcs.FSCD.2022.12.

Claude Stolze, Marino Miculan, and Pietro Di Gianantonio. Apr. 2023. "Composable Partial Multiparty Session Types for Open Systems". *Software and Systems Modeling*, 22, 2, (Apr. 2023), 473–494. DOI: 10.1007/s10270-022-01040-x.

Kaku Takeuchi, Kohei Honda, and Makoto Kubo. 1994. "An Interaction-Based Language and Its Typing System". In: *PARLE'94. Parallel Architectures and Languages Europe* (Lecture Notes in Computer Science) 10201. 6th International

PARLE Conference (Athens, Greece, July 4–8, 1994). Ed. by Costas Halatsis, Dimitrios Maritsas, George Philokyprou, and Sergios Theodoridis. Springer-Verlag Berlin Heidelberg, Berlin, 398–413. ISBN: 978-3-540-48477-6. DOI: 10.1007/3-540-581 84-7_118.

The Go Project. Feb. 6, 2024. *The Go Programming Language Specification*. Language version go1.22. (Feb. 6, 2024). Retrieved Mar. 5, 2024 from https://perma.cc/436Q-ZL6J.

Peter Thiemann and Vasco T. Vasconcelos. Dec. 2019. "Label-Dependent Session Types". *Proceedings of the ACM on Programming Languages*, 4, POPL, (Dec. 2019), 67. DOI: 10.1145/3371135.

Bernardo Toninho, Luís Caires, and Frank Pfenning. 2011. "Dependent Session Types Via Intuitionistic Linear Type Theory". In: *PPDP'11*. 13th International ACM SIGPLAN Symposium on Principles and Practices of Declarative Programming. PPDP'11 (Odense, Denmark, July 20–22, 2011). Association for Computing Machinery, Inc., New York, New York, 161–172. ISBN: 978-1-4503-0776-5. DOI: 10.1145/2003476.2003499.

Bernardo Toninho and Nobuko Yoshida. 2018. "Depending on Session-Typed Processes". In: *Foundations of Software Science and Computation Structures* (Lecture Notes in Computer Science) 10803. 21st International Conference, FOSSACS 2018 (Thessaloniki, Greece, Apr. 14–20, 2018). Ed. by Christel Baier and Ugo Dal Lago. European Joint Conferences on Theory and Practice of Software. SpringerOpen, Cham, Switzerland, 128–145. ISBN: 978-3-319-89366-2. DOI: 10.1007/978-3-319-8 9366-2_7.

Daniele Varacca and Nobuko Yoshida. Apr. 6, 2010. "Typed Event Structures and the Linear π-calculus". *Theoretical Computer Science*, 411, 19, (Apr. 6, 2010), 1949–1973. *Mathematical Foundations of Programming Semantics (MFPS 2006)*. DOI: 10.1016/j.tcs.2010.01.024.

Philip Wadler. Jan. 31, 2014. "Propositions As Sessions". *Journal of Functional Programming*, 24, 2-3, (Jan. 31, 2014), 384–418. DOI: 10.1017/s095679681400001x.