Deadlock and lock freedom in the linear π-calculus

Luca Padovani

▶ To cite this version:

Deadlock and lock freedom in the linear π-calculus

Luca Padovani – *Dipartimento di Informatica, Università di Torino, Italy*

Abstract—**We study two refinements of the linear π-calculus that ensure *deadlock freedom* (the absence of stable states with pending linear communications) and *lock freedom* (the eventual completion of pending linear communications). The main feature of both type systems is a new form of channel polymorphism that affects their accuracy in a significant way: they are the first of their kind that can deal with recursive processes communicating in cyclic network topologies.**

I. INTRODUCTION

The linear π-calculus [21] is a resource-aware model of communicating processes that distinguishes between linear and unlimited channels. Unlimited channels can be used without restrictions, while linear channels are meant to be used for *exactly one* communication. This intrinsic limitation is rewarded by several benefits, including specialized behavioral equivalences for reasoning about communication optimizations, the efficient implementation of linear channels, and the fact that communications on linear channels enjoy desirable properties such as determinism and confluence. The value of these benefits is amplified given that a significant fraction of channels in several actual systems happen to be linear.

From an operational standpoint, the linear π-calculus only guarantees that well-typed processes never communicate twice on the same linear channel. In practice, one may be interested in stronger guarantees, such as *deadlock freedom* [19] – the absence of stable states with pending communications on linear channels – or *lock freedom* [17], [22] – the possibility to complete pending communications on linear channels. A paradigmatic example of deadlock is illustrated by the process

$$a?(x).b!\langle x\rangle \mid b?(y).a!\langle y\rangle \tag{1}$$

where the left subprocess forwards on b the message x received from a, and the right subprocess forwards on a the message y received from b. The process (1) is well typed when a and b are linear channels, but none of the pending communications on a and b can complete because of the mutual dependencies between corresponding inputs and outputs. An example of lock which is not a deadlock is illustrated by the process

$$c!\langle a\rangle \mid *c?(x).c!\langle x\rangle \mid a!\langle 1984\rangle \tag{2}$$

where one occurrence of the linear channel a is repeatedly sent over the unlimited channel c but never used for receiving the 1984 message. Note that (2), unlike (1), reduces forever, but the communication pending on a cannot be completed.

In this work we propose two refinements of the linear π-calculus such that well-typed processes are (dead)lock free. The techniques we put forward have been inspired by previous ideas presented in [17]–[19], [22], except that by narrowing the focus on the linear π-calculus we obtain type systems that are not just technically simpler but also more accurate in establishing (dead)lock freedom of relevant processes.

In a nutshell, we consider channel types of the form $p^\iota[t]_m^n$ where the *polarity* p (either input ? or output !), the type of the message (t), and the *multiplicity* (unlimited ω or linear 1), are just like in the linear π-calculus. In addition, we annotate channel types with two numbers, a *priority* n and m *tickets*.

<mark>Priorities</mark> enforce an order on the use of channels: when a process owns two or more channels at the same time, channels with higher priority must be used *before* channels with lower priority (beware that we adopt the convention that "higher priority" means "smaller number"). In particular, an input $u?(x).P$ is well typed if u has higher priority than all the channels occurring in P and an output $u!\langle v\rangle$ is well typed if u has higher priority than v. This mechanism makes (1) ill typed: it is not possible to assign two priorities h and k to a and b, for the structure of the process requires the simultaneous satisfiability of the two constraints $h < k$ and $k < h$.

<mark>Tickets</mark> limit the number of travels that channels can do: $u!\langle v\rangle$ is well typed if v has at least one ticket; each time a channel is sent as a message, one ticket is removed from its type; a channel with no tickets cannot travel and must be used directly for performing a communication. This mechanism makes (2) ill typed, because a is sent on c infinitely many times and so it would need infinitely many tickets.

The technique described thus far prevents deadlocks (if we just consider the constraints on priorities) and locks (if we also consider the constraints on tickets). Unfortunately, it fails to type most recursive processes. For example, the usual encoding of the factorial below where, for convenience, we have annotated linear names with their priority, is ill typed:

$$*fact?(x, y^h). \tag{3}$$
$$\quad \text{if } x = 0 \text{ then } y^h!\langle 1\rangle$$
$$\quad \text{else } (\nu a^k)(fact!\langle x - 1, a^k\rangle \mid a^k?(z).y^h!\langle x \times z\rangle)$$

Since the newly created channel a is used in the same position as y in the recursive invocation of $fact$, we are led into thinking that a and y should have the same type hence the same priority $h = k$. This clashes with the input on a that blocks the output on y, requiring $k < h$. We see a symmetric phenomenon in

$$*stream?(x, y^h).(\nu a^k)(y^h!\langle x, a^k\rangle \mid stream!\langle x + 1, a^k\rangle) \tag{4}$$

which generates a stream of integers. Here too a and y are used in the same position and should have the same priority $h = k$, but the output $y!\langle x, a\rangle$ also requires $h < k$.

We need some way to overcome these difficulties. Let us digress for a moment from $fact$ and $stream$ and consider

$$F_1 \stackrel{\text{def}}{=} c?(x^h).y^k!\langle x^h\rangle \qquad \text{and} \qquad F_2 \stackrel{\text{def}}{=} c?(x^h, y^k).y^k!\langle x^h\rangle$$

both of which forward a linear channel x received from c to the linear channel y. In both cases it must be $k < h$, but there is a key difference between F_1 and F_2: the priority of x in F_1 has a fixed lower bound $k + 1$ because y is free, while the priority of x in F_2 is arbitrary, provided that $k < h$, because y is bound. Rephrased in technical jargon, c is *monomorphic* in F_1 and *polymorphic* in F_2 with respect to priorities. The fact that a channel like c is monomorphic or polymorphic depends on the presence or absence of free linear channels in the continuation that follows the input on c. Normally this information is not inferrable solely from the type of c, but it turns out that we can easily approximate it in the linear π-calculus, where replicated processes cannot have free linear channels because it is not known how many times they will run. So, unlimited channels that are *always* used for replicated inputs (*i.e.*, *replicated channels* [21]), are polymorphic. This is a very convenient circumstance, because replicated channels are the primary mechanism for implementing recursion.

Indeed, going back to *fact* and *stream*, we see that they are replicated channels, that is they are polymorphic. This means that the mismatches between the priorities h of y and k of a in (3) and (4) can be compensated by this form of polymorphism and these processes can be declared well typed.

The interplay between recursion and polymorphism leads to a technical problem, though. Recall that in (4) there are two occurrences of a, one in $stream!\langle x + 1, a\rangle$ having the same type as y but lower priority and another in $y!\langle x, a\rangle$ with a similar type, but opposite (input) polarity. Overall we realize that the type t of y should satisfy the equations

$$t = ![\texttt{int} \times t_1]^1 \qquad t_1 = ?[\texttt{int} \times t_2]^2$$
$$t_2 = ?[\texttt{int} \times t_3]^3 \qquad (5)$$
$$\vdots$$

where t_1 would be the type of a in $y!\langle x, a\rangle$ (for the sake of readability we omit tickets). The problem is that t is not a regular type and therefore cannot be finitely represented by means of the well-known μ notation [9]. To recover a finite representation for t, we "relativize" priorities. In particular, we postulate that t should satisfy the equations

$$t = ![\texttt{int} \times s]^1 \qquad s = ?[\texttt{int} \times s]^1$$

and that messages sent on a channel of type t must have type $\texttt{int} \times s$ (as written in t) but with priorities shifted up by 1 (the priority of t). Below we illustrate such shifting, which turns relative priorities (within channel types) into absolute ones:

Note that $\texttt{int}$ is not affected by the shifting because it is not a linear type and that overall the channel types in this sequence

are the same t_1, t_2, ... that we initially guessed in (5), except that now t and s are regular hence finitely representable.

Summary of contributions. We strengthen the notion of *linearity* in the linear π-calculus by defining two type systems ensuring the absence of deadlocks and locks involving linear channels, namely that linear channels are used *exactly* once as opposed to *at most* once. We exploit the features of the linear π-calculus to devise a form of channel polymorphism that allows us to deal with recursive processes that interleave actions on different linear channels also in cyclic network topologies. Such configurations are common in the implementation of parallel algorithms and session-based networks, and cannot be dealt with existing type systems for the generic π-calculus [17]–[19].

Outline. In Section II we quickly review syntax and semantics of the π-calculus and give formal definitions of deadlock and lock freedom. The type systems are described in Section III and illustrated on several examples that highlight the features of our approach. Section IV contains a more detailed and technical comparison with related work and particularly with [17]–[19]. Section V concludes and hints at ongoing and future developments. *Additional technical material, extended examples, and proofs of the results can be found in the appendix, which is not formally part of the submission.*

II. LANGUAGE

We use integer numbers h, k, ..., m, n, ..., a countable set of *variables* x, y, ..., and a countable set of *channels* a, b, ...; *names* u, v, ... are either channels or variables. *Expressions* e, ... and *processes* P, Q, ... are defined by:

$$e ::= n \mid u \mid (e, e) \mid \texttt{inl } e \mid \texttt{inr } e$$
$$P ::= \mathbf{0} \mid u?(x).P \mid *u?(x).P \mid u!\langle e\rangle \mid (P \mid Q) \mid (\nu a)P \mid$$
$$\texttt{let } x, y = e \texttt{ in } P \mid \texttt{case } e \; \{i \; x_i \Rightarrow P_i\}_{i=\texttt{inl},\texttt{inr}}$$

Expressions are integers, names, pairs of expressions, or expressions injected using either $\texttt{inl}$ or $\texttt{inr}$. *Values* v, w, ... are expressions without variables. Processes are the usual terms of the π-calculus enriched with pattern matching for pairs and injected values. In particular, $\texttt{let } x, y = e \texttt{ in } P$ deconstructs the value of e, which must be a pair, and continues as P where x and y have been respectively replaced by the first and second component of the pair. A process $\texttt{case } e \; \{i \; x_i \Rightarrow P_i\}_{i=\texttt{inl},\texttt{inr}}$ evaluates e, which must result into a value $(i \; \mathsf{v})$ where $i \in \{\texttt{inl}, \texttt{inr}\}$, and continues as P_i where x_i has been replaced by v. Binders are as expected, in particular $\texttt{let } x, y = e \texttt{ in } P$, binds both x and y in P, and $\texttt{case } e \; \{i \; x_i \Rightarrow P_i\}_{i=\texttt{inl},\texttt{inr}}$ binds x_i in P_i. The notions of *free* and *bound names* of a process P, respectively denoted by $\mathsf{fn}(P)$ and $\mathsf{bn}(P)$, are defined consequently. In the following we write $(\nu\tilde{a})$ for sequences of restrictions and we use polyadic inputs $u?(x_1, \ldots, x_n)$ as an abbreviation for monadic input of (possibly nested) pairs followed by (possibly nested) pair deconstructions using $\texttt{let}$'s. For example, $u?(x, y).P$ abbreviates $u?(z).\texttt{let } x, y = z \texttt{ in } P$ for some fresh z.

The operational semantics is defined by a combination of *structural congruence* $\equiv$ and a *reduction relation* $\rightarrow$. The former is the same as in the π-calculus, except that we do *not* include the law $*P \equiv P \mid *P$ because we treat replicated inputs in a special way. Reduction is defined by the rules below

$$a!\langle \mathsf{v}\rangle \mid a?(x).P \rightarrow P\{\mathsf{v}/x\}$$
$$a!\langle \mathsf{v}\rangle \mid *a?(x).P \rightarrow P\{\mathsf{v}/x\} \mid *a?(x).P$$
$$\texttt{let } x,y = \mathsf{v},\mathsf{w} \texttt{ in } P \rightarrow P\{\mathsf{v},\mathsf{w}/x,y\}$$
$$\texttt{case } k \ \mathsf{v} \ \{i \ x_i \Rightarrow P_i\}_{i=\texttt{inl,inr}} \rightarrow P_k\{\mathsf{v}/x_k\}$$

and closed by *reduction contexts* $\mathscr{C} ::= [\,] \mid (\mathscr{C} \mid P) \mid (\nu a)\mathscr{C}$ and structural congruence. The rules are unremarkable and $P\{\mathsf{v}/x\}$ is the usual capture-avoiding substitution of v in place of the free occurrences of x in P. We write $\rightarrow^*$ for the reflexive, transitive closure of $\rightarrow$ and $P \nrightarrow$ if there is no Q such that $P \rightarrow Q$. We say that P is *stable* if $P \nrightarrow$.

To formulate deadlock and lock freedom, we need a few predicates that describe the pending communications of a process P with respect to some channel a:

$$\begin{aligned}
\mathsf{in}(a,P) &\overset{\text{def}}{\Longleftrightarrow} P \equiv \mathscr{C}[a?(x).Q] \wedge a \notin \mathsf{bn}(\mathscr{C}) \\
*\mathsf{in}(a,P) &\overset{\text{def}}{\Longleftrightarrow} P \equiv \mathscr{C}[*a?(x).Q] \wedge a \notin \mathsf{bn}(\mathscr{C}) \\
\mathsf{out}(a,P) &\overset{\text{def}}{\Longleftrightarrow} P \equiv \mathscr{C}[a!\langle e\rangle] \wedge a \notin \mathsf{bn}(\mathscr{C}) \\
\mathsf{sync}(a,P) &\overset{\text{def}}{\Longleftrightarrow} (\mathsf{in}(a,P) \vee *\mathsf{in}(a,P)) \wedge \mathsf{out}(a,P) \\
\mathsf{wait}(a,P) &\overset{\text{def}}{\Longleftrightarrow} (\mathsf{in}(a,P) \vee \mathsf{out}(a,P)) \wedge \neg\mathsf{sync}(a,P)
\end{aligned}$$

In words, $\mathsf{in}(a,P)$ holds if there is a sub-process Q within P that is waiting for a message from a (it is similar to the *live* predicate in [4]). Note that, by definition of reduction context, the input is not guarded by other actions. The condition $a \notin \mathsf{bn}(\mathscr{C})$ means that a occurs free in P. The predicates $\mathsf{out}(a,P)$ and $*\mathsf{in}(a,P)$ are similar, but they regard outputs and replicated inputs, respectively. Whenever one of $\mathsf{in}(a,P)$ or $\mathsf{out}(a,P)$ holds, there is some pending input/output operation (we do not interpret $*\mathsf{in}(a,P)$ as a pending communication, because we do not require that a replicated input process should run infinitely often). This discussion explains the $\mathsf{sync}(a,P)$ and $\mathsf{wait}(a,P)$ predicates: the first one denotes the fact that there are pending input/output operations on a, but a synchronization on a is *immediately* possible; the second one denotes the fact that there is a pending output or a pending non-replicated input on a, but no immediate synchronization on a is possible. Note that pending outputs trigger the wait predicate even when they regard unlimited channels. This is because such outputs may carry messages containing linear channels and in order to guarantee deadlock freedom we must be sure that *all* outputs, including those on unlimited channels, can be completed and the messages delivered. For this reason we treat unlimited channels as replicated channels and enforce input receptiveness for them. On the contrary, replicated inputs, which can be though of as persistently available servers, do not trigger the wait predicate because, by definition, they are not allowed to have free linear channels (the type system will enforce this property).

We define deadlock freedom as the absence of stable states with pending communications and lock freedom as the ability to eventually complete any pending communication. Formally:

Definition II.1 (deadlock [19] and lock freedom [22])**.**
1) We say that P is *deadlock free* if for every Q such that $P \rightarrow^* (\nu\tilde{a})Q \nrightarrow$ we have $\neg\mathsf{wait}(a,Q)$ for every a.
2) We say that P is *lock free* if for every Q such that $P \rightarrow^* (\nu\tilde{a})Q$ and $\mathsf{wait}(a,Q)$ there exists R such that $Q \rightarrow^* R$ and $\mathsf{sync}(a,R)$.

For example, both $(\nu a)a!\langle 1984\rangle$ and $(\nu a)a?(x).P$ are deadlocks, whereas $(\nu a)*a?(x).P$ is lock free. In principle, we should also require that $\texttt{let}$ and $\texttt{case}$ processes are always able to reduce and that expressions in top-level outputs are always evaluated. These properties are already guaranteed (for closed processes) by conventional type systems for the π-calculus, so we keep Definition II.1 focused on pending communications. Note that lock freedom implies deadlock freedom, but not viceversa (process (2) is deadlock free but not lock free).

Example II.2. Many parallel algorithms (Jacobi and Gauss-Seidel, leader election, vertex coloring, just to mention a few) use batteries of processes that iteratively communicate with their neighbors. To maximize parallelism, communication is *full-duplex*, that is processes simultaneously send messages to each other. For instance, the process

$$Node_{\mathrm{A}} \overset{\text{def}}{=} *c_{\mathrm{A}}?(x,y).(\nu a)(x!\langle a\rangle \mid y?(z).c_{\mathrm{A}}!\langle a,z\rangle)$$

uses a channel x for sending messages to, and another channel y for receiving messages from, a neighbor process. Each message sent on x carries a payload (omitted) as well as a continuation channel a with which $Node_{\mathrm{A}}$ communicates with its neighbor at the next iteration. Symmetrically, each message received from y contains the neighbor's payload (omitted) and continuation z. The use of fresh continuations at each iteration makes sure that messages are received in the desired order. An alternative modeling using *half-duplex* communication is

$$Node_{\mathrm{B}} \overset{\text{def}}{=} *c_{\mathrm{B}}?(x,y).y?(z).(\nu a)(x!\langle a\rangle \mid c_{\mathrm{B}}!\langle a,z\rangle)$$

where the process first waits for the message from its neighbor, and only then sends its own information. It may then be relevant to understand whether a given configuration consisting of mixed nodes, some of type A and others of type B, is lock free. Below we represent three of them

$$\begin{aligned}
L_1(\mathrm{X}) &\overset{\text{def}}{=} Node_{\mathrm{A}} \mid Node_{\mathrm{B}} \mid c_{\mathrm{X}}!\langle e,e\rangle \\
L_2(\mathrm{X},\mathrm{Y}) &\overset{\text{def}}{=} Node_{\mathrm{A}} \mid Node_{\mathrm{B}} \mid c_{\mathrm{X}}!\langle e,f\rangle \mid c_{\mathrm{Y}}!\langle f,e\rangle \\
L_3(\mathrm{X},\mathrm{Y},\mathrm{Z}) &\overset{\text{def}}{=} Node_{\mathrm{A}} \mid Node_{\mathrm{B}} \mid c_{\mathrm{X}}!\langle e,f\rangle \mid c_{\mathrm{Y}}!\langle g,e\rangle \mid c_{\mathrm{Z}}!\langle f,g\rangle
\end{aligned}$$

where each of X, Y, and Z can be either A and B to yield a different configuration. For example, we have

$$\begin{aligned}
L_2(\mathrm{B},\mathrm{B}) \rightarrow^* \ &Node_{\mathrm{A}} \mid Node_{\mathrm{B}} \mid f?(z).(\nu a)(e!\langle a\rangle \mid c_{\mathrm{B}}!\langle a,z\rangle) \\
&\mid e?(z).(\nu b)(f!\langle a\rangle \mid c_{\mathrm{B}}!\langle b,z\rangle) \nrightarrow
\end{aligned}$$

and the final configuration satisfies the in predicate but not out for both e and f, namely $L_2(\mathrm{B},\mathrm{B})$ is *not* lock free. We will see later that a given configuration L_i is lock free if and only if at least one of its nodes is of type A. $\blacksquare$

III. TYPE SYSTEM

Notation. *Polarities* p, q, ... are subsets of $\{?, !\}$. We abbreviate $\{?\}$ with $?$, $\{!\}$ with $!$, and $\{?, !\}$ with $\#$ and we say that $\emptyset$ and $\#$ are *even polarities*. *Multiplicities* ι, ... are either 1 or ω. We also use a countable set of *type variables* α, ... *Types* t, s, ... are defined by the grammar below:

$$ t ::= \texttt{int} \mid \alpha \mid t \times s \mid t \oplus s \mid p^{\iota}[t]_m^n \mid \mu\alpha.t $$

The types $\texttt{int}$, $t \times s$, and $t \oplus s$ respectively denote integers, pairs inhabited by values (v, w) where v has type t and w has type s, and the disjoint sum of t and s inhabited by values $(\texttt{inl }\mathsf{v})$ when v has type t or $(\texttt{inr }\mathsf{w})$ when w has type s. The type $p^{\iota}[t]_m^n$ denotes a channel to be used with polarity p and multiplicity ι for exchanging messages of type t. The polarity determines the operations allowed on the channel: $\emptyset$ means none, $?$ means input, $!$ means output, and $\#$ means both. The multiplicity determines how many times the channel can or must be used: 1 means that the channel must be used exactly once (for each element in p), while ω means that the channel can be used any number of times. The numbers n and m are respectively the *priority* and the *tickets* of the channel. As we have anticipated in Section I, priorities are used for imposing an order on the usage of channels, so that an action on a channel with low priority cannot block another action on a channel with higher priority. Recall that smaller numbers denote higher priorities, so for example an input on a channel with priority 4 can block an input on a channel with priority 5, but not viceversa. The tickets of a channel limit the number of times the channel can be sent in a message. For example, a channel with 3 tickets can be sent *at most* three times in a message, but it can be used at any time for an input/output operation according to its polarity. A channel with no tickets cannot be sent in a message. Tickets are always non-negative. We omit priority and tickets when $\iota = \omega$ and the multiplicity when it is 1. Also, we often omit tickets when there is none. We use type variables and μ's for building recursive types. Notions of free and bound type variables are as expected.

Contractiveness. For simplicity, we forbid non-contractive types in which recursion variables are not guarded by a channel type. For example, $\mu\alpha.\alpha$ and $\mu\alpha.(\alpha \times \alpha)$ are illegal while $\mu\alpha.p^{\iota}[\alpha]_m^n$ is allowed. In this way, we will be able to define functions by induction on the structure of types, when such functions do not recur within channel types. Contractiveness can be weakened without compromising the results that follow, but some functions on types must be defined more carefully. The formal definition of contractiveness can be found in the appendix. We identify two types modulo renaming of bound type variables and if they have the same infinite unfolding, that is if they denote the same (regular) tree [9]. In particular, we have $\mu\alpha.t = t\{\mu\alpha.t/\alpha\}$.

Priorities and tickets. The priority of a channel gives a measure to the urgency with which the channel must be used: the higher the priority is, the sooner the channel must be used. This measure can be extended from channels to arbitrary types, the idea being that the priority of a type is the highest of the priorities of the channel types occurring therein. To compute the priority of a type, we define an auxiliary function $|\cdot|$ such that $|t|$ is an element of $\mathbb{Z} \cup \{\bot, \top\}$ where $\bot < n < \top$ for every $n \in \mathbb{Z}$. The function is defined inductively thus:

$$ |t| \overset{\text{def}}{=} \begin{cases} \bot & \text{if } t = p^{\omega}[s] \text{ and } ? \in p \\ n & \text{if } t = p[s]_m^n \text{ and } p \neq \emptyset \\ \min\{|t_1|, |t_2|\} & \text{if } t = t_1 \times t_2 \text{ or } t = t_1 \oplus t_2 \\ \top & \text{otherwise} \end{cases} \quad (6) $$

Intuitively, unlimited channels with input polarity have the highest priority $\bot$ (first equation) because their use cannot be postponed by any means to guarantee input receptiveness. Numbers, unlimited channels with only output polarity, and linear channels with empty polarity have the lowest priority $\top$ (last equation) because they do not affect (dead)lock freedom in any way. Linear channels with pending operations must be used according to their priority (second equation). The priority of compound types is the highest (numerically speaking, the minimum) of the priorities of the component types (third equation). For instance, $|![t]_1^3 \times ?[s]_4^2| = \min\{|![t]_1^3|, |?[s]_4^2|\} = \min\{3, 2\} = 2$.

We say that a (value with) type t is *unlimited* if $|t| = \top$, that it is *linear* if $|t| \in \mathbb{Z}$, that it is *relevant* if $|t| = \bot$.

We define another auxiliary function $\$_k^h$ to *shift* priorities and tickets: $\$_k^h t$ has the same structure of t, except that all priorities/tickets in the topmost linear channel types of t have been transposed by h and k respectively. Formally:

$$ \$_k^h t \overset{\text{def}}{=} \begin{cases} p[s]_{m+k}^{n+h} & \text{if } t = p[s]_m^n \text{ and } p \neq \emptyset \\ (\$_k^h t_1) \times (\$_k^h t_2) & \text{if } t = t_1 \times t_2 \\ (\$_k^h t_1) \oplus (\$_k^h t_2) & \text{if } t = t_1 \oplus t_2 \\ t & \text{otherwise} \end{cases} \quad (7) $$

As an example, we have $\$_1^1 (\texttt{int} \times ?[\texttt{int}]^3) = (\$_1^1 \texttt{int}) \times (\$_1^1 ?[\texttt{int}]^3) = \texttt{int} \times ?[\texttt{int}]_1^4$. Note that positive/negative shifting respectively corresponds to decreasing/increasing the priority. We write $\$^n$ for $\$_0^n$.

A key property used in the proof of subject reduction is that shifting behaves like the identity on unlimited types:

Lemma III.1. *If t is unlimited, then $\$_k^h t = t$ for any h, k.*

Monomorphic and polymorphic channels. A distinguishing feature of our type systems is that message types are "relative" to the priority of the channels on which they travel. We use shifting to compute the absolute type of messages that can be sent on a channel. For example, a channel whose type is $![s]^h$ accepts messages of type $\$^h s$, that is the type s (in the type t of the channel) where all top-level priorities have been shifted by h (the priority of the channel). As we have seen in Section I, this feature allows us to work with finite representations of recursive types having infinitely many different priorities. For example, a recursive type that satisfies the equation $t = ![t]^1$ denotes a linear channel on which it is possible to send a message of type $![t]^2$. In turn, on such a channel it is possible to send a message of type $![t]^3$, and so on and so forth.

Shifting is also used to implement polymorphism. For example, an unlimited (*i.e.*, polymorphic) channel of type

4

$!^\omega[s]$ accepts messages of type $\$^h s$ *for any* h. So, for linear (*i.e.*, monomorphic) channels, the shifting is fixed and determined by the priority of the channel, while for unlimited (*i.e.*, polymorphic) channels, the shifting is arbitrary.

We are aware that this way of realizing polymorphism, as opposed to *e.g.* using "priority variables", is sometimes more restrictive than necessary. For example, a channel of type $!^\omega[?[\texttt{int}]^1 \times !\,[\texttt{int}]^0]$ accepts messages of type

$$\$^h\,(?[\texttt{int}]^1 \times !\,[\texttt{int}]^0) = ?[\texttt{int}]^{h+1} \times !\,[\texttt{int}]^h$$

for every h, but not messages of type $?[\texttt{int}]^{h+2} \times !\,[\texttt{int}]^h$ where the first component of the pair has been shifted by $h+1$ and the second component of the pair only by h. A message of the latter type could be safely sent on the unlimited channel c appearing in the process F_2 of Section I. The type systems allow more flexible forms of polymorphism (see Section V).

Type environments. We check that processes are well typed in <mark>type environments</mark> $\Gamma, \dots,$ which are finite maps from names to types written $u_1 : t_1, \dots, u_n : t_n$. We write $\mathsf{dom}(\Gamma)$ for the *domain* of Γ, namely the set of names for which there is a binding in Γ, and Γ, Γ' for the union of Γ and Γ' when $\mathsf{dom}(\Gamma) \cap \mathsf{dom}(\Gamma') = \emptyset$. In general we need a more flexible way of composing type environments, taking into account the linearity and relevance of types and the fact that we can split channel types by distributing different polarities to different processes. Following [21], we define a partial composition operator $+$ between types, thus:

$$
\begin{aligned}
t + t &= t && \text{if } t \text{ is unlimited} \\
p^\omega[t] + q^\omega[t] &= (p \cup q)^\omega[t] && \quad\quad (8) \\
p[s]_h^n + q[s]_k^n &= (p \cup q)[s]_{h+k}^n && \text{if } p \cap q = \emptyset
\end{aligned}
$$

Informally, unlimited types compose with themselves without restrictions. The composition of two unlimited/relevant channel types has the union of their polarities. Two linear channel types can be composed only if they have the same priority and disjoint polarities, and the composition has the union of their polarities and the sum of their tickets. We extend the partial operator $+$ to type environments, thus:

$$
\begin{aligned}
\Gamma + \Gamma' &\stackrel{\text{def}}{=} \Gamma, \Gamma' && \text{if } \mathsf{dom}(\Gamma) \cap \mathsf{dom}(\Gamma') = \emptyset \\
(\Gamma, u : t) + (\Gamma', u : s) &\stackrel{\text{def}}{=} (\Gamma + \Gamma'), u : t + s &&
\end{aligned}
$$
$$(9)$$

Note that $\Gamma + \Gamma'$ is undefined if there is $u \in \mathsf{dom}(\Gamma) \cap \mathsf{dom}(\Gamma')$ such that $\Gamma(u) + \Gamma'(u)$ is undefined and that $\mathsf{dom}(\Gamma + \Gamma') = \mathsf{dom}(\Gamma) \cup \mathsf{dom}(\Gamma')$. We write $\mathsf{un}(\Gamma)$ if all the types in the range of Γ are unlimited. We let $|\Gamma|$ denote the maximum priority of the types in the range of Γ, that is $|\Gamma| = \min\{|\Gamma(u)| \mid u \in \mathsf{dom}(\Gamma)\}$.

Typing rules. The typing rules for expressions and processes are presented in Table 1. A judgment $\Gamma \vdash e : t$ denotes that e is well typed and has type t in Γ and a judgment $\Gamma \vdash_k P$ denotes that P is well typed in Γ. The only difference between the type systems for deadlock and lock freedom is the "cost" k of travels for linear channels, which is 0 for deadlock freedom and 1 for lock freedom. This parameter affects only the rules for outputs, and is always propagated unchanged by all rules.

TABLE 1: Typing rules for expressions and processes.

Expressions

$$
\frac{\mathsf{un}(\Gamma)}{\Gamma \vdash n : \texttt{int}} \text{ [T-CONST]} \qquad
\frac{\mathsf{un}(\Gamma)}{\Gamma, u : t \vdash u : t} \text{ [T-NAME]}
$$

$$
\frac{\Gamma \vdash e : t \quad \Gamma' \vdash e' : s}{\Gamma + \Gamma' \vdash (e, e') : t \times s} \text{ [T-PAIR]} \qquad
\frac{\Gamma \vdash e : t}{\Gamma \vdash \texttt{inl}\ e : t \oplus s} \text{ [T-INL]} \qquad
\frac{\Gamma \vdash e : s}{\Gamma \vdash \texttt{inr}\ e : t \oplus s} \text{ [T-INR]}
$$

Processes

$$
\frac{\mathsf{un}(\Gamma)}{\Gamma \vdash_k \mathbf{0}} \text{ [T-IDLE]} \qquad
\frac{\Gamma, x : \$^n\, t \vdash_k P \quad n < |\Gamma|}{\Gamma + u : ?[t]_m^n \vdash_k u?(x).P} \text{ [T-IN]} \qquad
\frac{\Gamma \vdash e : \$_k^n\, t \quad 0 < |t|}{\Gamma + u : !\,[t]_m^n \vdash_k u!\langle e \rangle} \text{ [T-OUT]}
$$

$$
\frac{\Gamma, x : t \vdash_k P \quad \mathsf{un}(\Gamma)}{\Gamma + u : ?^\omega[t] \vdash_k *u?(x).P} \text{ [T-IN*]} \qquad
\frac{\Gamma \vdash e : \$_k^n\, t \quad \perp < |t|}{\Gamma + u : !^\omega[t] \vdash_k u!\langle e \rangle} \text{ [T-OUT*]}
$$

$$
\frac{\Gamma, a : p^t[t]_m^n \vdash_k P \quad p \text{ even}}{\Gamma \vdash_k (\nu a)P} \text{ [T-NEW]} \qquad
\frac{\Gamma \vdash e : t \times s \quad \Gamma', x : t, y : s \vdash_k P}{\Gamma + \Gamma' \vdash_k \texttt{let}\ x, y = e\ \texttt{in}\ P} \text{ [T-LET]}
$$

$$
\frac{\Gamma_i \vdash_k P_i^{\,(i=1,2)}}{\Gamma_1 + \Gamma_2 \vdash_k P_1 \mid P_2} \text{ [T-PAR]} \qquad
\frac{\Gamma \vdash e : t \oplus s \quad \Gamma', x : t \vdash_k P \quad \Gamma', y : s \vdash_k Q}{\Gamma + \Gamma' \vdash_k \texttt{case}\ e\ \{\texttt{inl}\ x \Rightarrow P, \texttt{inr}\ y \Rightarrow Q\}} \text{ [T-CASE]}
$$

Rules for expressions as well as [T-IDLE], [T-PAR], [T-LET], and [T-CASE] are standard for the linear π-calculus, with the usual splitting of environments and the requirement that unused environments in axioms must be unlimited.

Rule [T-IN] is used for typing linear inputs $u?(x).P$, where u must have type $?[t]_m^n$. The continuation P is typed in an environment where the input polarity of u has been removed and the received message x has been added. Note that the type of x is not just t, but t shifted by n, consistently with the relative interpretation of priorities of message types. The tickets are irrelevant since u is used for an input operation, not as the content of a message. The condition $n < |\Gamma|$ verifies that the input on u does not block operations on other channels with equal or higher priority. In particular, Γ cannot contain relevant channels. Below are some typical examples of ill-typed processes that violate this condition:

- $a : ?[\texttt{int}]^1, b : !\,[\texttt{int}]^0 \vdash_k a?(x).b!\langle x \rangle$ is not derivable because $1 \not< 0$: the input on a blocks the output on b, but b has higher priority than a;
- $a : \#[\texttt{int}]^h \vdash_k a?(x).a!\langle x \rangle$ is a degenerate case of the previous example, where the input on a blocks the very output that should synchronize with it. Note that this process is well typed in the traditional linear π-calculus because $\#[\texttt{int}] = ?[\texttt{int}] + !\,[\texttt{int}]$;
- $a : ?[\texttt{int}]^h, c : ?^\omega[\texttt{int}] \vdash_k a?(x).*c?(y)$ is not derivable because $|?^\omega[\texttt{int}]| = \perp$. To guarantee input receptiveness, replicated inputs cannot be guarded by other operations.

Rule [T-OUT] is used for typing linear outputs on channels of type $!\,[t]_m^n$. The type of the message e must be t (as specified in the type of the channel u) shifted by n again since t is

relative to the priority of u. The tickets are shifted by 1, but only for lock freedom ($k = 1$). The meaning is that, by sending e as a message, one ticket from every linear channel type in the type of e is consumed. This prevents the degenerate phenomenon of infinite delegation that we have seen in (2). The condition $0 < |t|$ verifies that the priority of the message is lower (*i.e.*, numerically greater) than that of the channel on which it travels. Below are three examples:

- The judgment $a : \,![?[\texttt{int}]_0^1]_0^2, b : \,?[\texttt{int}]_1^3 \vdash_1 a!\langle b \rangle$ is derivable because $?[\texttt{int}]_1^3 = \$_1^2\,?[\texttt{int}]_0^1$. Note in particular that the channel to be sent on a must have no tickets, which is in fact what happens to b after 1 ticket is consumed from its type before it travels on a.

- The judgment $a : \,?[![\texttt{int}]_0^2]_0^0, b : \,![![\texttt{int}]_0^1]_0^1 \vdash_1 a?(x).b!\langle x \rangle$ is not derivable, because x received from a has no tickets so it cannot travel on b.

- Let $t = \mu\alpha.?[\alpha]^0$ and observe that $\#[t]^1 = \,![t]^1 + \,?[t]^1$. The judgment $a : \#[t]^1 \vdash_0 a!\langle a \rangle$ is not derivable, despite the message a has the "right" type $?[t]^1 = \$^1\,t$, because the condition $0 < |t| = 0$ is not satisfied. According to Definition II.1, a process like $(\nu a)a!\langle a \rangle$ is deadlock. (there are systems in which such process actually reduces, but in doing so it generates memory leaks [2]).

Rule [T-IN*] is used for typing replicated inputs $*u?(x).P$. This rules differs from [T-IN] in three important ways. First of all, u must be an unlimited channel with input polarity. Second, the residual environment Γ must be unlimited, because the continuation P can be spawned an arbitrary number of times. Third, it may be the case that $u \in \text{dom}(\Gamma)$, because $?^\omega[t] + \,!^\omega[t] = \#^\omega[t]$ according to (8) and $!^\omega[t]$ is unlimited. This means that replicated input processes may invoke themselves, as we have seen in many examples.

Rule [T-OUT*] is used for outputs on unlimited channels. There are two key differences with respect to [T-OUT]. First, the condition $\bot < |t|$, where t is the type of x, implies that only relevant names cannot be communicated, but a process like $a!\langle a \rangle$ is well typed when a is unlimited, for example by giving the inner a type $\mu\alpha.!^\omega[\alpha]$. Second, the type of the message need not match exactly the type t in the channel, but its priority can be shifted by an arbitrary amount n. This is the technical realization of polymorphism. In particular, each distinct output on u can shift the type of the message by a different amount, therefore allowing polymorphic recursion. We will see this feature at work Example III.6.

Rule [T-NEW] is used for restricting new (linear and unlimited) channels. Only channels with an even polarity can be restricted: either the channel comes with full polarity $\#$, or with no polarity $\emptyset$ (this is only useful for subject reduction).

We write $\vdash_{\text{DF}} P$ if $\emptyset \vdash_0 P$ and $\vdash_{\text{LF}} P$ if $\emptyset \vdash_1 P$ and all priorities in the derivation are non-negative. Intuitively, the type system for deadlock freedom ignores tickets because infinite delegations (like in (2)) are not deadlocks, while the type system for lock freedom requires a well-founded order on priorities (hence the restriction to natural numbers) and an upper bound on delegations (hence the importance of tickets).

Properties. The relative interpretation of priorities makes it possible to shift whole derivations and preserve typing. More precisely, let $\$^n\,\Gamma$ be the environment obtained by shifting all the types in the range of Γ by n, that is $(\$^n\,\Gamma)(u) = \$^n\,\Gamma(u)$ for every $u \in \text{dom}(\Gamma)$. We have:

Lemma III.2 (shifting). *The following properties hold:*
1) *If $\Gamma \vdash e : t$, then $\$^n\,\Gamma \vdash e : \$^n\,t$.*
2) *If $\Gamma \vdash_0 P$, then $\$^n\,\Gamma \vdash_0 P$.*
3) *If $\Gamma \vdash_1 P$ and $n \geq 0$, then $\$^n\,\Gamma \vdash_1 P$.*

Note that the converse of 3) does not hold. For example, the derivation for $u : \,![\texttt{int}]_0^1 \vdash_1 (\nu a)(a!\langle 1984 \rangle \mid a?(x).u!\langle x \rangle)$ relies on the fact that u has priority 1, for otherwise it would not be possible to prefix $u!\langle x \rangle$ with an input operation on a.

Typing is preserved by reductions and the type environment may change as a consequence of communications on linear channels, just like in the linear π-calculus [21]:

Lemma III.3 (subject reduction). *If $\Gamma \vdash_k P$ and $P \to P'$, then $\Gamma' \vdash_k P'$ for some Γ'.*

The proof is essentially standard, except for the interesting case concerning the communication on unlimited (*i.e.*, polymorphic) channels sketched below. Consider the derivation

$$\frac{\dfrac{\Gamma \vdash \mathsf{v} : \$^n\,t}{a : \,!^\omega[t], \Gamma \vdash_k a!\langle \mathsf{v} \rangle} \qquad \dfrac{\Gamma', x : t \vdash_k P \quad \text{un}(\Gamma')}{a : \,?^\omega[t], \Gamma' \vdash_k *a?(x).P}}{a : \#^\omega[t], \Gamma + \Gamma' \vdash_k a!\langle \mathsf{v} \rangle \mid *a?(x).P}$$

and observe that the actual type $\$^n\,t$ of the message v being sent does *not* match exactly the type t expected by the receiver, but is shifted by some arbitrary amount n. In order to type the reduct $P\{\mathsf{v}/x\}$, we proceed like this. First of all we shift the derivation of $\Gamma', x : t \vdash_k P$ using Lemma III.2, and obtain:

$$\$^n\,\Gamma', x : \$^n\,t \vdash_k P$$

Then, we observe that Γ' is unlimited because it is the type environment used for typing a replicated process. This means that $\$^n\,\Gamma' = \Gamma'$, since shifting is the identity on unlimited types (Lemma III.1). Therefore we deduce:

$$\Gamma', x : \$^n\,t \vdash_k P$$

Now, the actual and expected types of the message v coincide and we know from the initial derivation that $\Gamma + \Gamma'$ is defined, so we can apply the conventional substitution lemma for the linear π-calculus [21] and conclude:

$$\Gamma + \Gamma' \vdash_k P\{\mathsf{v}/x\}$$

Since the type systems refine the one in [21], all the properties of the linear π-calculus (partial confluence, linear usage of channels, etc.) still hold. The added value is that the refined type systems guarantee deadlock/lock freedom.

Theorem III.4 (soundness). *The following properties hold:*
1) *If $\vdash_{\text{DF}} P$, then P is deadlock free.*
2) *If $\vdash_{\text{LF}} P$, then P is lock free.*

6

TABLE 2: Typing derivation for $Node_B$.

$$\cfrac{\cfrac{\cfrac{a : \$_2^n s \vdash a : \$_2^n s \qquad 0 < m}{x : t, a : \$_2^n s \vdash_1 x!\langle a\rangle}\ {\scriptstyle[\text{T-OUT}]} \qquad \cfrac{z : \$_1^m s, a : \$_1^m t \vdash (a,z) : \$_1^m (t \times s)}{c_B : !^\omega[t \times s], z : \$_1^m s, a : \$_1^m t \vdash_1 c_B!\langle a, z\rangle}\ {\scriptstyle[\text{T-OUT*}]}}{\cfrac{c_B : !^\omega[t \times s], x : t, z : \$_1^m s, a : \#[\$_1^0 s]_3^{m+n} \vdash_1 x!\langle a\rangle \mid c_B!\langle a, z\rangle}{\cfrac{c_B : !^\omega[t \times s], x : t, z : \$_1^m s \vdash_1 (\nu a)(x!\langle a\rangle \mid c_B!\langle a, z\rangle) \qquad\qquad m < n}{\cfrac{c_B : !^\omega[t \times s], x : t, y : s \vdash_1 y?(z).(\nu a)(x!\langle a\rangle \mid c_B!\langle a, z\rangle)}{c_B : \#^\omega[t \times s] \vdash_1 *c_B?(x,y).y?(z).(\nu a)(x!\langle a\rangle \mid c_B!\langle a, z\rangle)}\ {\scriptstyle[\text{T-IN*}]}}\ {\scriptstyle[\text{T-IN}]}}\ {\scriptstyle[\text{T-NEW}]}}\ {\scriptstyle[\text{T-PAR}]}$$

The proof of 1) is by contradiction, for the existence of a well-typed, stable process with pending input or output operations on linear channels would imply the existence of an infinite chain of channels with different priorities, contradicting the fact that there are only finitely many distinct channels in the process. The proof of 2) does an induction on a *measure* that includes, among other information, priority and tickets of the linear channel for which we are completing the communication (see Definition II.1(2)). The induction is well founded if so is the domain of priorities, whence the restriction to natural numbers for the priorities in the type system for lock freedom. The details can be found in the appendix.

We conclude this section with a few examples showcasing the key features of our type systems. In particular, the recursive processes in Examples III.5, III.7, and III.8 are representative of those configurations that are typeable thanks to our form of channel polymorphism, but not in similar type systems [17]–[19] (see Section IV for a detailed discussion).

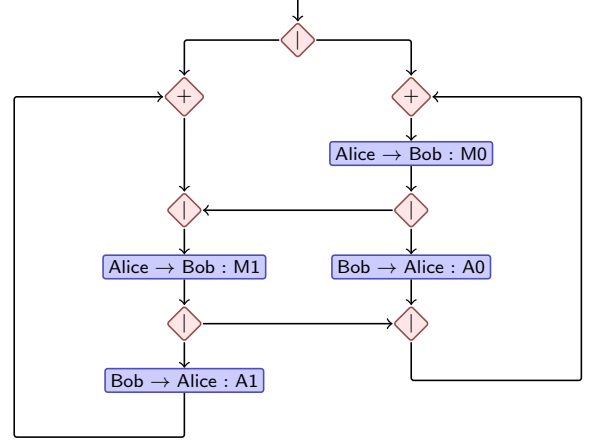**Example III.5.** Regarding the processes in Example II.2, let

$$t \stackrel{\text{def}}{=} ![s']_0^n \quad \text{and} \quad s \stackrel{\text{def}}{=} ?[s']_0^m \quad \text{and} \quad s' \stackrel{\text{def}}{=} \mu\alpha.?[\alpha]_1^m$$

and observe that $s' = \$_1^0 s = ?[\$_1^0 s]_1^m$. For $Node_B$ we obtain the typing derivation shown in Table 2 if and only if $0 < m < n$. An analogous derivation is obtained for $Node_A$ if and only if $0 < n, m$. We deduce that $L_1(A)$, $L_2(A, A)$, $L_2(A, B)$, and $L_2(B, A)$ are all well typed. Since these constraints on priorities are the most general ones, we also deduce that there are no derivations for $L_1(B)$ and $L_2(B, B)$, suggesting that these configurations are locked. ■

Example III.6. Below is the encoding of the function that computes the n-th number in the sequence of Fibonacci. Instead of showing the typing derivation, which is conventional for the linear π-calculus, we just annotate each linear name occurring in the term with its priority and leave it to the reader to verify that such annotations are consistent with the constraints expressed in the typing rules for deadlock freedom:

```
*fibo?(x, y⁰).
   if x ≤ 1 then y⁰!⟨x⟩
   else (νa⁻¹)(νb⁻²)(fibo!⟨x − 1, a⁻¹⟩ | fibo!⟨x − 2, b⁻²⟩ |
                    b⁻²?(z₂).a⁻¹?(z₁).y⁰!⟨z₁ + z₂⟩)
```

Note the use of negative priorities associated with the fresh continuation channels a and b, due to the fact that the inputs

Fig. 1: Alternating Bit Protocol [12].

on these channels block the output on the continuation y. Observe also the role of polymorphic recursion in typing this process: the two recursive invocations of *fibo* are applied to continuation channels a and b with different priorities. ■

Example III.7. The following variation of the process $Node_A$ gives an example of dynamic lock-free system (beware that in this and the next example, subscripts in names denote tickets; we use colors to ease readability):

$$Ring \stackrel{\text{def}}{=} *c?(x_0^0, y_0^0).$$
$$(\nu a_3^1)\big(x_0^0?(z_1^1).(\nu b_2^1)(c!\langle z_1^1, b_1^1\rangle \mid c!\langle b_1^1, a_1^1\rangle) \mid y_0^0!\langle a_2^1\rangle\big)$$

Each message on c spawns two duplicates connected by a new channel b. As a consequence, the process $(\nu e)(Ring \mid c!\langle e, e\rangle)$ grows as a ring of processes, each sending a message to its neighbor, and the ring doubles its diameter each time a round of communications is completed. ■

Example III.8. Figure 1 depicts a variant of the Alternating Bit Protocol [12] in which Alice sends alternated messages M0 and M1 to Bob, and Bob answers with corresponding acknowledgments A0 and A1. The diagram describes communications (the rectangles) between the two parties and their dependencies (the arcs), combined with merging points (+ nodes), forks (| nodes with two outgoing arcs) and joins (| nodes with two incoming arcs). The particularity of this modeling of the protocol is that Alice concurrently waits for

both A0 and A1 before sending the next corresponding Mi, but always alternating between M0 and M1 messages.

We model Bob as the process below, which is parametric in two channels x and y from which Bob respectively receives messages M0 and M1 along with the continuations on which he sends the corresponding Ai (we omit the actual payloads and focus on channels and their continuations):

$$\mathsf{Bob} \stackrel{\text{def}}{=} *bob?(x^0, y^2).$$
$$(\nu a_3^3 b_3^5)\big(x^0?(\bar{x}^1).(\bar{x}^1!\langle a_2^3 \rangle \mid y^2?(\bar{y}^3).(\bar{y}^3!\langle b_2^5 \rangle \mid bob!\langle a_1^3, b_1^5 \rangle))\big)$$

We model Alice as two parallel threads Alice0 and Alice1, each handling messages and acknowledgments with matching number, so that Alice0 corresponds to the rightmost vertical flow in Figure 1, and Alice1 to the leftmost one. To respect the protocol and prevent that two messages with the same index are sent in a row, Alice uses a third channel z (unknown to Bob) with which she synchronizes the two threads:

$$\mathsf{Alice0} \stackrel{\text{def}}{=} *alice0?(x^0, z^1).$$
$$(\nu a_1^1 c_1^2)\big(x^0!\langle a_1^1 \rangle \mid z^1!\langle c_1^2 \rangle \mid a^1?(\bar{x}_1^3).c^2?(\bar{z}_1^4).alice0!\langle \bar{x}_1^3, \bar{z}_1^4 \rangle\big)$$

$$\mathsf{Alice1} \stackrel{\text{def}}{=} *alice1?(y^2, z^1).$$
$$(\nu b_1^3 c_3^4)\big(z^1?(\bar{z}^2).(y^2!\langle b_1^3 \rangle \mid \bar{z}^2!\langle c_2^4 \rangle \mid b^3?(\bar{y}_1^5).alice1!\langle \bar{y}_1^5, c_1^4 \rangle)\big)$$

The whole protocol is modeled as the term $(\nu ab)(\mathsf{Bob} \mid \mathsf{Alice0} \mid \mathsf{Alice1} \mid bob!\langle a, b \rangle \mid (\nu c)(alice0!\langle a, c \rangle \mid alice1!\langle b, c \rangle))$. Note the interleaving of possibly blocking actions on different channels in both Alice and Bob. ∎

IV. RELATED WORK

The articles [17]–[19], [22] have been a source of inspiration for our work so we devote most of this section to them. The discussion is quite technical and space constraints force us to assume some familiarity with these works.

Deadlock freedom. [19] defines a type system for ensuring deadlock freedom in the π-calculus. Channels types are pairs $[t]/U$ where t is the type of messages and U – called *usage* – is a term of a process algebra that includes prefixing and parallel composition and that specifies how the channel is accessed. For example, a channel of type $[\mathtt{int}]/(!_n^m.?_k^h \mid ?_m^n.!_h^k)$ can be used concurrently by two processes, one that first sends and then receives an integer and the other that does the opposite. Actions like $!_n^m$ are annotated with an *obligation* m and a *capability* n to prevent mutual dependencies between channels: inputs on a channel with capability n can only block operations on other channels whose obligation is larger than n. The special value ∞ is allowed for denoting various forms of unlimited channels. There is a significant overlapping between [19] and our work: usages can describe linear channels as a particular case; linearized channels (those that can be accessed multiple times, but only sequentially) can be encoded using linear channels and continuation passing; priorities have been directly inspired by obligations and capabilities and are used similarly. Let us now analyze the main differences.

First we observe that having two quantities (obligation/capability) instead of a single one (priority) may improve

the accuracy of the analysis in some cases. For example,

$$*c?(x, y).x?(z).y?(z) \mid *c?(x, y).y?(z).x?(z)$$

is ill typed in our type system but is typeable in [19] by giving both x and y obligation 1 and capability 0. The null capability means that, assuming that outputs on x and y are available in the environment immediately (capability 0), then either of these processes guarantees that the both inputs will be performed after at most one reduction (obligation 1).

The main distinguishing features of our approach, however, are polymorphism and the fact that priorities have a relative interpretation in message types. To illustrate the effect of these features on the accuracy of our type systems, we revisit Example II.2 in the setting of [19] assuming to extend usages with recursion (recursive usages are admitted in [19], but only for type reconstruction purposes). The following modeling of $L_2(\mathsf{A}, \mathsf{A})$ uses no explicit continuation passing and therefore is the most favorable to the type system in [19]:

$$*c?(x, y).(x!\langle 3 \rangle \mid y?(z).c!\langle x, y \rangle) \mid c!\langle e, f \rangle \mid c!\langle f, e \rangle \quad (10)$$

The process is tentatively typed using the assignments

$$e : [\mathtt{int}]/(\mu\alpha.!_n^m.\alpha \mid \mu\alpha.?_m^n.\alpha)$$
$$f : [\mathtt{int}]/(\mu\alpha.?_k^h.\alpha \mid \mu\alpha.!_h^k.\alpha)$$

where the fact that obligations and capabilities are swapped in complementary actions of the usages of e and f is necessary to satisfy the *reliability* condition [19] which corresponds, in our setting, to the property that channel types with opposite polarities and same priority can be combined together (last equation of (8)). The structure of (10) requires the simultaneous satisfiability of $k < m$ (there is an input on f with capability k that blocks e whose topmost obligation is m) and of $m < k$ (there is an input on e with capability m that blocks f whose topmost obligation is k), therefore (10) cannot be proven deadlock free with the type discipline in [19]. The problem is that obligations and capabilities, unlike priorities, are absolute values associated with actions: e and f always have the same obligation and capability regardless of the number of unfoldings of their usages. By contrast, in Example III.5 the channels e and f and their continuations are assigned a numerically increasing sequence of priorities. There are two features in [19] that overcome partially the limitations due to absolute obligations and capabilities. The first one is that the behavioral nature of usages allows the typing of recursive processes involving one channel. For example, the following variant of $L_2(\mathsf{A}, \mathsf{B})$

$$*c?(y).(y!\langle\rangle \mid c!\langle y \rangle) \mid *d?(y).y?(z).d!\langle y \rangle \mid c!\langle e \rangle \mid d!\langle e \rangle$$

is well typed in [19]. The point is that in [19] an input process $u?(x).P$ is well typed if the capability of u is smaller than the obligation of all the channels in P *except u itself*. The second feature stems from the observation that the strict order $<$ used for comparing the capability of a channel u and the obligation of another channel v can be safely weakened to $\leq$ if u has been created later than v. This property, denoted by the relation $u \prec v$ in [19], can be inferred from the syntactic structure

of processes and makes it possible to deal with processes like (3), where the input on the newer channel a is allowed to block operations on the older channel y despite the fact that a and y have the same type (hence the same obligations and capabilities). As noted in [19], the very same feature would also work using the opposite order $\succ$, but one would have to choose between either $\prec$ or $\succ$, since using their union would be unsound. The opposite order $\succ$ however arises naturally with continuation passing, as we have seen in (4) and in most of the examples throughout the paper.

In summary, the techniques based on absolute obligations/capabilities have difficulties handling systems where two or more recursive processes interleave actions on two or more channels. The use of the $\prec$ relation alleviates in part these limitations, when processes have a function-like behavior.

Usages provide more precise information on unlimited channels. For example, the lock-free modeling of the dining philosophers in [17] cannot be captured by our type system because it uses non-replicated unlimited channels for representing shared resources (the forks) accessed by competing processes (the philosophers).

Lock freedom. Type systems for lock freedom have been presented in [17], [18]. These works rely on essentially the same types and usages already discussed for deadlock freedom, and so they are subject to the same problems that arise in (10). In particular, all the processes discussed in Examples III.5, III.7, III.8, which are representative configurations mixing recursion and cyclic network topologies, are ill typed in [17]–[19].

The partial order $\prec$ is unsound and cannot be used to improve the accuracy of lock-free analysis in [17], [18]. This is reflected also in our type system for lock freedom, where negative priorities are banned. Consequently, processes such as (3) or the one in Example III.6 can be proved deadlock free but not lock free. Lock-free typability of (3) (restricted to natural numbers) and other recursive-like processes is recovered in [22], where it is observed that lock freedom can be characterized as the conjunction of deadlock freedom and *termination*. Such characterization turns out to be more accurate on recursive processes with function-like behaviors, partly because the type system for deadlock freedom imposes fewer constraints, partly thanks to the accuracy of the techniques for proving termination. In fact, the hybrid approach [22] is parametric on the specific techniques for verifying deadlock freedom or termination. This means that it suffers from the same shortcomings that we have discussed earlier for processes like (10) when instantiated with the deadlock free analysis of [19], but also that it should be applicable in our setting, in which case it would benefit from the better accuracy of our type system for deadlock freedom. By the way, the limit we impose on the number of times a channel can be sent in a message is akin to ensuring a form of termination (in fact, the eventual use of the channel for performing a communication) so in a sense our type system for lock freedom can already be seen as an instance of the hybrid framework described in [22].

Session types. Session types [14], [15] are behavioral types akin to sequential usages in which each single input/output action is annotated with the type of a message. For example, the session type $!{[}\texttt{int}{]}.?{[}\texttt{bool}{]}$ denotes a channel on which it is possible to first sent an integer, and then receive a boolean. *Duality* relates session types associated with corresponding endpoints of a binary session and generalizes the composition operator between linear types (8). For example, the session type above is dual of $?{[}\texttt{int}{]}.!{[}\texttt{bool}{]}$. Duality ensures communication safety and also lock freedom, but the latter property is guaranteed only *within* sessions, not when (blocking) actions from different sessions are interleaved with each other. In some works [4], [26] a plain session type discipline ensures deadlock freedom also when sessions are interleaved, but only because the syntax of (well-typed) processes prevents the modeling of cyclic network topologies.

Global types. A global type [5], [12], [13], [16] describes the *interactions* between participants of a session as opposed to the *actions* that participants perform on the channel of the session. For example, the interactions between the three processes that compose the system $L_3(\texttt{A}, \texttt{B}, \texttt{B})$ in Example III.5 can be described by the recursive global type G_3 that satisfies the equation $G_3 = \texttt{A} \to \texttt{B}.\texttt{B} \to \texttt{C}.\texttt{C} \to \texttt{A}.G_3$ which gives names $\texttt{A}$, $\texttt{B}$, and $\texttt{C}$ to the three *participants* and specifies the order of their synchronizations. In general, global types also allow the specification of the type of the exchanged messages and include operators other than prefixing, such as choice and parallel composition. Global types that satisfy some well-formedness conditions are *projectable* into session types and processes that conform to these projections are guaranteed to be type safe, to implement the protocol described by the global type, and to be lock free. Therefore, global types are an effective way of extending the lock freedom property from binary to n-ary sessions.

One downside of global types is that they may make it harder to build and maintain systems *compositionally* (although some modularity and compositionality aspects have recently been addressed in [11] and [23]). For instance, $L_2(\texttt{A}, \texttt{B})$ and $L_3(\texttt{A}, \texttt{B}, \texttt{B})$ in Example II.2 are assembled in a modular way using the same constituents, but the global types $G_2 = \texttt{A} \to \texttt{B}.\texttt{B} \to \texttt{A}.G_2$ and G_3 above that describe the interactions in these systems are quite different. It may also be harder to describe "unstructured" interactions from a global perspective. For instance, it takes the most sophisticated global type language available to date [12] to describe accurately the dependencies between Alice's threads in Example III.8. For these reasons, global types are ideal for describing delimited interactions within sessions, but they cannot dispense completely from the need of interleaving different sessions, in which case they are unable to prevent (dead)locks.

These observations have motivated the study of mixed approaches [1], [7], [8] that keep track of the order in which different sessions interleave with the purpose of detecting and flagging mutual dependencies between sessions that could lead to (dead)locks. These mixed approaches turn out to be quite restrictive (*e.g.*, recursive processes can access one session only [8]) and require heavyweight and multi-layer type systems.

V. Conclusions

Type systems that enforce (dead)lock freedom of communicating processes inevitably rely on trade-offs, for such properties are undecidable [3] even under severe restrictions of the communication model (*e.g.*, for half-duplex communication with three or more processes [6]). These type systems essentially follow one of two approaches: some [17]–[19] favor compositionality and apply to generic process algebras, but have difficulties in dealing with recursive processes and cyclic network topologies; others can handle such configurations, but only within the scope of (multiparty) sessions and assuming that the communication protocol is explicitly described in advance by a global type. In the context of the linear π-calculus, our type systems conjugate the advantages of both approaches and do so with minimal machinery. Indeed, all previous type systems for (dead)lock freedom rely on rich behavioral types (usages [17]–[20], [22], session types [24], global types [1], [7], [8], [12], [13], [16], conversation types [25]), and when they are able to deal with recursive processes it is because they take advantage of such richness (see Section IV), almost implying the necessity of rich types to ensure properties as strong as (dead)lock freedom. We have shown that this is not the case, at least for the linear π-calculus.

The linear π-calculus is an adequate model for many concrete communicating systems. In particular, it is the "assembly language" that lies beneath binary sessions [10], [20] and in the appendix we show that many multiparty sessions too can be compiled into the linear π-calculus. So, one can establish (dead)lock freedom of a session-based process (also in presence of session interleaving) by compiling it into the linear π-calculus and by typing the resulting process using our type systems. When this approach is unfeasible – not all multiparty sessions can be encoded using linear channels – or undesirable – the structure of sessions may serve other important purposes – the technique described in this paper can still be used for reasoning on inter-session (dead)lock freedom: just like we annotate *channel types* with priorities and tickets, it is possible to annotate the single *actions* within session and global types in the same way. This experiment was already attempted in [24], [25], but in these works the annotations are "absolute", much like obligations/capabilities, and so they suffer from issues analogous to those described in Section IV.

The form of channel polymorphism that lies at the core of our type systems allows the shifting of priorities with maps of the shape $z \mapsto z + n$. This prevents, for example, typing an output $a!\langle x^0, y^1 \rangle$ if a has type $!^\omega[?[\texttt{int}]^1 \times \,![\texttt{int}]^3]$, because a expects x and y to have priorities at distance 2, while they are at distance 1. In fact, the type system for deadlock freedom can be generalized to rational numbers for priorities and *positive affine maps* for transforming them. In the example above, the map $z \mapsto \frac{1}{2}z - \frac{1}{2}$ turns the priorities 1 and 3 in the type of a into 0 and 1, namely the priorities of x and y. The type system for lock freedom is slightly less flexible, because its soundness proof requires priorities to be part of a well-founded set.

Finally, as witnessed by the examples in this paper, it is un-realistic to assume that the programmer is capable of guessing the correct values of priorities and tickets of linear channels in general. We are finalizing a complete type reconstruction algorithm for our type systems. Intuitively, the algorithm uses variables for representing unknown priorities and tickets and accumulates constraints between such variables according to the rules of the type systems (Table 1). These constraints can always be represented as an integer programming problem, which is solvable using known methods.

References

[1] L. Bettini, M. Coppo, L. D'Antoni, M. De Luca, M. Dezani-Ciancaglini, and N. Yoshida. Global progress in dynamically interleaved multiparty sessions. In *CONCUR'08*, LNCS 5201, pages 418–433, 2008.

[2] V. Bono and L. Padovani. Typing Copyless Message Passing. *Log. Meth. in Comp. Sci.*, 8:1–50, 2012.

[3] D. Brand and P. Zafiropulo. On communicating finite-state machines. *Journal of the ACM*, 30(2):323–342, 1983.

[4] L. Caires and F. Pfenning. Session types as intuitionistic linear propositions. In *CONCUR'10*, LNCS 6269, pages 222–236, 2010.

[5] G. Castagna, M. Dezani-Ciancaglini, and L. Padovani. On Global Types and Multi-Party Sessions. *Log. Meth. in Comp. Sci.*, 8:1–45, 2012.

[6] G. Cécé and A. Finkel. Verification of programs with half-duplex communication. *Inf. Comput.*, 202(2):166–190, Nov. 2005.

[7] M. Coppo, M. Dezani-Ciancaglini, L. Padovani, and N. Yoshida. Inference of Global Progress Properties for Dynamically Interleaved Multiparty Sessions. In *COORDINATION'13*, LNCS 7890, pages 45–59. Springer, 2013.

[8] M. Coppo, M. Dezani-Ciancaglini, N. Yoshida, and L. Padovani. Global progress for dynamically interleaved multiparty sessions. *Math. Struct. in Comp. Sci.*, to appear.

[9] B. Courcelle. Fundamental properties of infinite trees. *Theor. Comp. Sci.*, 25:95–169, 1983.

[10] O. Dardha, E. Giachino, and D. Sangiorgi. Session types revisited. In *PPDP'12*, pages 139–150. ACM, 2012.

[11] R. Demangeon and K. Honda. Nested protocols in session types. In *CONCUR'12*, LNCS 7454, pages 272–286. Springer, 2012.

[12] P.-M. Deniélou and N. Yoshida. Multiparty session types meet communicating automata. In *ESOP'12*, LNCS 7211, pages 194–213. Springer, 2012.

[13] P.-M. Deniélou, N. Yoshida, A. Bejleri, and R. Hu. Parameterised multiparty session types. *Log. Meth. in Comp. Sci.*, 8(4), 2012.

[14] K. Honda. Types for dyadic interaction. In *CONCUR'93*, LNCS 715, pages 509–523. Springer, 1993.

[15] K. Honda, V. T. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, LNCS 1381, pages 122–138. Springer, 1998.

[16] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *POPL'08*, pages 273–284. ACM, 2008.

[17] N. Kobayashi. A type system for lock-free processes. *Inf. and Comp.*, 177(2):122–159, 2002.

[18] N. Kobayashi. Type-based information flow analysis for the pi-calculus. *Acta Inf.*, 42(4-5):291–347, 2005.

[19] N. Kobayashi. A new type system for deadlock-free processes. In *CONCUR'06*, LNCS 4137, pages 233–247. Springer, 2006.

[20] N. Kobayashi. Type systems for concurrent programs. Technical report, Tohoku University, 2007. Short version appeared in 10th Anniversary Colloquium of UNU/IIST, 2002.

[21] N. Kobayashi, B. C. Pierce, and D. N. Turner. Linearity and the pi-calculus. *ACM Trans. Program. Lang. Syst.*, 21(5):914–947, 1999.

[22] N. Kobayashi and D. Sangiorgi. A hybrid type system for lock-freedom of mobile processes. *ACM Trans. Program. Lang. Syst.*, 32(5), 2010.

[23] F. Montesi and N. Yoshida. Compositional choreographies. In *CONCUR'13*, LNCS 8052, pages 425–439. Springer, 2013.

[24] L. Padovani. From Lock Freedom to Progress Using Session Types. In *PLACES'13*, EPTCS 137, pages 3–19, 2013.

[25] H. T. Vieira and V. T. Vasconcelos. Typing progress in communication-centred systems. In *COORDINATION'13*, LNCS 7890, pages 236–250. Springer, 2013.

[26] P. Wadler. Propositions as sessions. In *ICFP'12*, pages 273–286. ACM, 2012.

Although the reduction relation defined in Section II is adequate to explain the operational semantics of the language and the deadlock and lock freedom properties, it lacks some important information that is instrumental for the proof of subject reduction. In particular, it is necessary to keep track of the *channel* on which a synchronization occurs, so that we can suitably update the type of that channel in case it is linear (recall that, by definition, only one synchronization is supposed to occur on a linear channel). To this aim, we define a *labeled variant* of the reduction relation which keeps track of this information. Labels $act, \ldots$ are either channels or the special symbol τ indicating an unobservable reduction. The labeled reduction relation is inductively defined by the rules below:

$$[\text{LR-COMM}]$$
$$a!\langle v \rangle \mid a?(x).P \xrightarrow{a} P\{v/x\}$$

$$[\text{LR-COMM*}]$$
$$a!\langle v \rangle \mid *a?(x).P \xrightarrow{a} P\{v/x\} \mid *a?(x).P$$

$$[\text{LR-LET}]$$
$$\texttt{let } x, y = (v, w) \texttt{ in } P \xrightarrow{\tau} P\{v, w/x, y\}$$

$$[\text{LR-CASE}]$$
$$\texttt{case } (k \; v) \; \{i \; x_i \Rightarrow P_i\}_{i=\texttt{inl,inr}} \xrightarrow{\tau} P_k\{v/x_k\}$$

$$[\text{LR-PAR 1}] \qquad \frac{P \xrightarrow{act} P'}{P \mid Q \xrightarrow{act} P' \mid Q}$$

$$[\text{LR-PAR 2}] \qquad \frac{Q \xrightarrow{act} Q'}{P \mid Q \xrightarrow{act} P \mid Q'}$$

$$[\text{LR-NEW 1}] \qquad \frac{P \xrightarrow{act} Q \qquad act \neq a}{(\nu a)P \xrightarrow{act} (\nu a)Q}$$

$$[\text{LR-NEW 2}] \qquad \frac{P \xrightarrow{a} Q}{(\nu a)P \xrightarrow{\tau} (\nu a)Q}$$

$$[\text{LR-STRUCT}] \qquad \frac{P \equiv P' \qquad P' \xrightarrow{act} Q' \qquad Q' \equiv Q}{P \xrightarrow{act} Q}$$

The two reduction semantics are perfectly equivalent:

Theorem A.1. *The predicates $P \to Q$ and $\exists act : P \xrightarrow{act} Q$ are equivalent.*

Proof. A straightforward induction on the derivation of the reduction. $\qquad\square$

A. Definitions

Below is the formal definition of contractive types: we say that a type is *contractive* if $\emptyset \vdash \texttt{con}(t)$ is inductively derivable by the rules

$$A \vdash \texttt{con}(\texttt{int}) \qquad A \setminus \{\alpha\} \vdash \texttt{con}(\alpha) \qquad \frac{\emptyset \vdash \texttt{con}(t)}{A \vdash \texttt{con}(p^\iota[t]_n^m)}$$

$$\frac{A \vdash \texttt{con}(t_i) \; {}^{(i=1,2)}}{A \vdash \texttt{con}(t_1 \times t_2)} \qquad \frac{A \vdash \texttt{con}(t_i) \; {}^{(i=1,2)}}{A \vdash \texttt{con}(t_1 \oplus t_2)} \qquad \frac{A \cup \{\alpha\} \vdash \texttt{con}(t)}{A \vdash \texttt{con}(\mu\alpha.t)}$$

where A is a set of type variables. Note that $\mu\alpha.p^\iota[\alpha]_n^m$ is contractive but neither $\mu\alpha.(\alpha \times \alpha)$ nor $\mu\alpha.(\alpha \oplus \alpha)$ are. The last one, in particular, seems preventing the definition of any recursive algebraic type. The definition of contractiveness we adopt is convenient for the rest of the technical development but it can be relaxed to allow α be guarded by *any* constructor, at the cost of definiting the function $|\cdot|$ coinductively (see (6)).

B. Basic properties

Below are two auxiliary results regarding the type systems, namely the shifting property and the substitution lemma for expressions and processes. Both results can be easily proved with an induction on the derivation of the judgment in the premise.

Lemma B.1 (Lemma III.2). *The following properties hold:*

1) *If $\Gamma \vdash e : t$, then $\$^n \Gamma \vdash e : \$^n t$.*
2) *If $\Gamma \vdash_0 P$, then $\$^n \Gamma \vdash_0 P$.*
3) *If $\Gamma \vdash_1 P$ and $n \geq 0$, then $\$^n \Gamma \vdash_1 P$.*

Lemma B.2 (substitution). *Let $\Gamma' \vdash v : t$ and $\Gamma + \Gamma'$ be defined. Then:*

1) *$\Gamma, x : t \vdash e : s$ implies $\Gamma + \Gamma' \vdash e\{v/x\} : s$.*
2) *$\Gamma, x : t \vdash_k P$ implies $\Gamma + \Gamma' \vdash_k P\{v/x\}$.*

We will make extensive use of Lemma III.2 in the proofs that follow, to the point that it is convenient to extend the type system presented in Table 1 with the (admissible) rule

$$\frac{[\text{T-SHIFT}] \quad \Gamma \vdash_k P \qquad k = 1 \Rightarrow n \geq 0}{\$^n \, \Gamma \vdash_k P}$$

C. Subject reduction

In order to formulate and prove the subject reduction result, we need to determine a precise correspondence between the type environment *before* a reduction and the one *after* the reduction. We do so by defining a reduction relation for type environments using the same labels as those in the labeled reduction for processes. Let $\xrightarrow{act}$ be the least relation defined by the rules

$$\Gamma, a : \#[t]_m^n \xrightarrow{a} \Gamma, a : \emptyset[t]_m^n \qquad \Gamma, a : \#^\omega[t] \xrightarrow{a} \Gamma, a : \#^\omega[t] \qquad \Gamma, a : \$_1^0 \, t \xrightarrow{\tau} \Gamma, a : t$$

The first rule consumes a linear channel a from the environment and is used when a synchronization on a occurs (recall that linear channels can be used for one synchronization only by definition). The middle rule leaves the environment unchanged when a synchronization on some unlimited channel a occurs. The last rule decreases consumes one ticket from the type of a and is used each time a is sent in a message. In the following we write $\xRightarrow{\tau}$ for the reflexive, transitive closure of $\xrightarrow{\tau}$ and $\xRightarrow{act}$ for the composition $\xRightarrow{\tau} \xrightarrow{act} \xRightarrow{\tau}$.

Lemma B.3. *Let* $\Gamma \vdash v : \$_1^0 \, t$. *Then there exists* Γ' *such that* $\Gamma \xRightarrow{\tau} \Gamma'$ *and* $\Gamma' \vdash v : t$.

Proof. By induction on the derivation of $\Gamma \vdash v : \$_1^0 \, t$. We only show two cases, correponding to the channel and pair expressions:

- $(v = a)$ Then $\Gamma = \Gamma'', a : p^\iota[s]_m^n$ and $un(\Gamma'')$. We distinguish two subcases:
  - $(\iota = 1)$ Then $\$_1^0 \, t = p[s]_m^n$ and $m > 0$ and $t = p[s]_{m-1}^n$ and we conclude by taking $\Gamma' = \Gamma'', a : t$ and observing that $\Gamma \xrightarrow{\tau} \Gamma''$.
  - $(\iota = \omega)$ Then $\$_1^0 \, t = t = p^\omega[s]$ and we conclude by taking $\Gamma' = \Gamma$.
- $(v = (v_1, v_2))$ Then $\Gamma = \Gamma_1 + \Gamma_2$ and $\$_1^0 \, t = (\$_1^0 \, t_1) \times (\$_1^0 \, t_2)$ and $\Gamma_i \vdash v_i : \$_1^0 \, t_i$ for $i = 1, 2$. By induction hypothesis there exist Γ_1' and Γ_2' such that $\Gamma_i \xRightarrow{\tau} \Gamma_i'$ and $\Gamma_i' \vdash v_i : t_i$ for $i = 1, 2$. Observe that the $\xrightarrow{\tau}$ relation between type environments can only consume the tickets of values. This means that $\Gamma_1' + \Gamma_2'$ is defined and moreover $\Gamma \xRightarrow{\tau} \Gamma_1' + \Gamma_2'$. We conclude by taking $\Gamma' = \Gamma_1' + \Gamma_2'$. $\square$

The next result is standard and indicates that typing is preserved by structural congruence. Note that the result would not hold, had we included the law $*P \equiv P \mid *P$.

Lemma B.4. *If* $\Gamma \vdash_k P$ *and* $P \equiv Q$, *then* $\Gamma \vdash_k Q$.

Proof. By induction on the derivation of $P \equiv Q$. $\square$

The following result states that well-typing of processes is preserved by reductions. Note though that the type environment may change.

Lemma B.5 (Lemma III.3). *If* $\Gamma \vdash_k P$ *and* $P \xrightarrow{act} P'$, *then* $\Gamma' \vdash_k P'$ *for some* Γ' *such that* $\Gamma \xRightarrow{act} \Gamma'$.

Proof. By induction on the derivation of $P \xrightarrow{act} P'$ and by cases on the last rule applied. We only show the most relevant cases.

$\boxed{[\text{LR-COMM}]}$ Then $P = a!\langle v \rangle \mid a?(x).Q \xrightarrow{a} Q\{v/x\} = P'$ and $act = a$. From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ where $\Gamma_1 \vdash_k a!\langle v \rangle$ and $\Gamma_2 \vdash_k a?(x).Q$. From [T-OUT] we deduce $\Gamma_1 = \Gamma_1' + a : ![t]_m^n$ and $\Gamma_1' \vdash v : \$_k^n \, t$. From [T-IN] we deduce $\Gamma_2 = \Gamma_2' + a : ?[t]_h^n$ and $\Gamma_2', x : \$^n \, t \vdash_k Q$. The fact that the types associated with the two a's are complementary follows from the definition of $\Gamma_1 + \Gamma_2$. By Lemma B.3 we deduce that there exists Γ_1'' such that $\Gamma_1'' \vdash v : \$^n \, t$ and $\Gamma_1' \xRightarrow{\tau} \Gamma_1''$. Note that $\Gamma_1'' + \Gamma_2'$ is defined. Then by Lemma B.2(2) we obtain $\Gamma_1'' + \Gamma_2' \vdash_k Q\{v/x\}$. Let $\Gamma' = \Gamma_1'' + \Gamma_2' + a : \emptyset[t]_{m+h}^n$. From [T-PAR] we deduce $\Gamma' \vdash_k P'$. We conclude by observing that $\Gamma \xRightarrow{act} \Gamma'$.

$\boxed{[\text{LR-COMM*}]}$ Then $P = a!\langle v \rangle \mid *a?(x).Q \xrightarrow{a} Q\{v/x\} \mid *a?(x).Q = P'$ and $act = a$. From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ where $\Gamma_1 \vdash_k a!\langle v \rangle$ and $\Gamma_2 \vdash_k *a?(x).Q$. From [T-OUT*] we deduce $\Gamma_1 = \Gamma_1' + a : !^\omega[t]$ and $\Gamma_1' \vdash e : \$_k^n \, t$. From [T-IN*] we deduce $\Gamma_2 = \Gamma_2' + a : ?^\omega[t]$ and $\Gamma_2', x : t \vdash_k Q$ and $un(\Gamma_2')$. By Lemma B.3 we deduce that there exists Γ_1'' such that $\Gamma_1'' \vdash v : \$^n \, t$ and $\Gamma_1' \xRightarrow{\tau} \Gamma_1''$. By Lemma III.2 we deduce $\$^n \, \Gamma_2', x : \$^n \, t \vdash_k Q$. From $un(\Gamma_2')$ and Lemma III.1 we deduce $\Gamma_2' = \$^n \, \Gamma_2'$, therefore $\Gamma_1' + \$^n \, \Gamma_2'$ is defined. By Lemma B.2(2) we deduce $\Gamma_1'' + \Gamma_2' \vdash_k Q\{v/x\}$. From [T-IN*] and [T-PAR] we obtain $\Gamma_1'' + \Gamma_2 \vdash_k Q\{v/x\} \mid *a?(x).Q$. We conclude by taking $\Gamma' = \Gamma_1'' + \Gamma_2$ and observing that $\Gamma \xRightarrow{act} \Gamma'$.

$\boxed{[\text{LR-LET}]}$ Then $P = \text{let } x, y = (\mathsf{v}, \mathsf{w}) \text{ in } Q$ and $P' = Q\{\mathsf{v}, \mathsf{w}/x, y\}$. From [T-LET] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash \mathsf{v} : t \times s$ and $\Gamma_2, x : t, y : s \vdash_k Q$. From [T-PAIR] we deduce $\Gamma_1 = \Gamma_{11} + \Gamma_{12}$ where $\Gamma_{11} \vdash \mathsf{v} : t$ and $\Gamma_{12} \vdash \mathsf{w} : s$. By Lemma B.2 we deduce $\Gamma \vdash_k Q\{\mathsf{v}, \mathsf{w}/x, y\}$ and we conclude by taking $\Gamma' = \Gamma$.

$\boxed{[\text{LR-CASE}]}$ Then $P = \text{case } (i\ \mathsf{v})\ \{\text{inl } x_1 \Rightarrow P_1, \text{inr } x_2 \Rightarrow P_2\} \to P_i\{\mathsf{v}/x_i\} = P'$ for some $i \in \{\text{inl}, \text{inr}\}$. From [T-CASE] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ where $\Gamma_1 \vdash (i\ \mathsf{v}) : t_1 \oplus t_2$ and $\Gamma_2, x_i : t_i \vdash_k P_i$. From either [T-INL] or [T-INR] we deduce $\Gamma_1 \vdash \mathsf{v} : t_i$. From Lemma B.2 we deduce $\Gamma \vdash_k P_i\{\mathsf{v}/x_i\}$ and we conclude by taking $\Gamma' = \Gamma$.

$\boxed{[\text{LR-NEW 1}]}$ Then $P = (\nu a)Q$ and $Q \xrightarrow{act} Q'$ and $P' = (\nu a)Q'$ and $act \neq a$. From the hypothesis $\Gamma \vdash_k P$ we deduce $\Gamma, a : p^\iota[t]_m^n \vdash_k Q$ where p is even. By induction hypothesis we derive that there exists Γ' and $h \leq m$ such that $\Gamma', a : p^\iota[t]_h^n \vdash_k Q'$ and $\Gamma \xrightarrow{act} \Gamma'$. We conclude with an application of [T-NEW]. $\qquad \square$

D. Soundness

Definition B.6 (balanced environment). We say that Γ is <mark>*balanced*</mark> if $\Gamma = \{a_i : p_i{}^{\iota_i}[t_i]_{m_i}^{n_i}\}_{i \in I}$ and p_i is even for every $i \in I$.

Lemma B.7. *Let $\emptyset \vdash_0 P$ and $P \nrightarrow$. Then $P \equiv (\nu \tilde{a})\left(\prod_{i \in I} *a_i?(x_i).P_i\right)$.*

Proof. From the hypothesis that P is well typed and does not reduce we know that $P \equiv (\nu \tilde{a})Q$ where

$$Q = \prod_{i \in I} *a_i?(x_i).P_i \mid \prod_{j \in J} b_j?(y_j).Q_j \mid \prod_{k \in K} c_k!\langle e_k \rangle$$

hence, from [T-NEW], we deduce $\Gamma \vdash_0 Q$ for some Γ that is balanced. We proceed showing that $J = K = \emptyset$ by contradiction. In particular, we show that if $J \cup K \neq \emptyset$, then it is possible to find an infinite chain of channels with strictly decreasing priorities. This is absurd since Q is finite, so there is a finite number of distinct channels in it.

We begin doing a few generic considerations about the channels occurring in these processes:

- all the a_i's are unlimited channels, because they are used for replicated inputs;
- all the b_j are linear channels, because they are used for linear inputs;
- all the c_k's are linear channels, because $P \nrightarrow$ whereas if they were unlimited they would be able to synchronize with one of the replicated inputs (recall that the type system guarantees input receptiveness);
- from [T-IN*] we know that none of the P_i's has free linear channels.

Now suppose $j \in J$ and $\Gamma(b_j) = ?[t]_m^n$. Since b_j is a linear channel, there must be another occurrence of b_j with output polarity. We observe that:

- b_j is not any of the b_l for $l \in J \setminus \{j\}$, for the occurrence of b_j we are looking for must have output polarity;
- b_j is not any of the c_k's, because $P \nrightarrow$ and all the e_k must be closed and therefore evaluated.

We analyze the remaining possibilities:

- ($b_j \in \mathsf{fn}(Q_l)$ for some $l \in J$) Then from [T-IN] we deduce $\Gamma(b_l) = ?[s]_{m'}^{n'}$ where $n' < n$.
- ($b_j \in \mathsf{fn}(e_k)$ for some $k \in K$) Then from [T-OUT] we deduce $\Gamma(c_k) = ![s]_{m'}^{n'}$ where $n' < n$.

Now suppose $k \in K$ and $\Gamma(c_k) = ![t]_m^n$. Since c_k is a linear channel, there must be another occurrence of c_k with input polarity. We observe that:

- c_k is not any of the b_j's, because $P \nrightarrow$.
- c_k is not any of the c_l for $l \in K \setminus \{j\}$, for the occurrence of c_k we are looking for must have input polarity.

We analyze the remaining possibilities:

- ($c_k \in \mathsf{fn}(Q_j)$ for some $j \in J$) Then from [T-IN] we deduce $\Gamma(b_j) = ?[s]_{m'}^{n'}$ where $n' < n$.
- ($c_k \in \mathsf{fn}(e_l)$ for some $l \in K$) Then from [T-OUT] we deduce $\Gamma(c_l) = ![s]_{m'}^{n'}$ where $n' < n$.

In conclusion, for every possibility that is not ruled out by simple considerations, from any channel in $\{b_j \mid j \in J\} \cup \{c_k \mid k \in K\}$ we have been able to find another channel in the same set with a strictly smaller priority. This contradicts the fact that the very same set must be finite. $\qquad \square$

In order to prove that the type system for lock freedom is sound we need to define the *measure* of a channel. Intuitively, the measure of a represents an upper bound to the (finite) number of reduction steps that must occur before a synchronization on a occurs. The measure of a channel depends on both its type and the positions in which it occurs in a process. We only consider processes that are well typed in a balanced environment.

We let $\mathsf{tickets}(t)$ denote the tickets of a type, which is computed similarly to the priority of a type, except that we consider m in the second equation of (6).

Definition B.8 (measure). Let $\Gamma \vdash_1 P$ where Γ is balanced and $a : t \in \Gamma$. Then $\mathsf{measure}(a, \Gamma, P) \stackrel{\text{def}}{=} (|t|, \mathsf{tickets}(t), [\![P]\!]_a)$ where $[\![P]\!]_a$ is inductively defined below, where we assume $\bot + 1 = \bot$:

$$\begin{aligned}
[\![a]\!]_a &= 0 \\
[\![u]\!]_a &= \bot && \text{if } a \neq u \\
[\![e_1, e_2]\!]_a &= [\![e_1]\!]_a \vee [\![e_2]\!]_a \\
[\![\texttt{inl } e]\!]_a = [\![\texttt{inr } e]\!]_a &= [\![e]\!]_a
\end{aligned}$$

$$\begin{aligned}
[\![\mathbf{0}]\!]_a &= \bot \\
[\![u?(x).P]\!]_a &= [\![u]\!]_a \vee ([\![P]\!]_a + 1) \\
[\![*P]\!]_a &= [\![P]\!]_a \\
[\![u!\langle e\rangle]\!]_a &= [\![u]\!]_a \vee [\![e]\!]_a \\
[\![P \mid Q]\!]_a &= [\![P]\!]_a \vee [\![Q]\!]_a \\
[\![\texttt{let } x, y = e \texttt{ in } P]\!]_a &= [\![e]\!]_a \vee [\![P]\!]_a \\
[\![\texttt{case } e \ \{i \ x_i \Rightarrow P_i\}_{i=\texttt{inl},\texttt{inr}}]\!]_a &= [\![e]\!]_a \vee [\![P_{\texttt{inl}}]\!]_a \vee [\![P_{\texttt{inr}}]\!]_a \\
[\![(\nu a)P]\!]_a &= \bot \\
[\![(\nu b)P]\!]_a &= [\![P]\!]_a && \text{if } a \neq b
\end{aligned}$$

Given $\Gamma \vdash_1 P$ where Γ is balanced and $a \in \mathsf{dom}(\Gamma)$, the measure of a is a triple consisting of the priority $|\Gamma(a)|$ of a, the tickets $\mathsf{tickets}(\Gamma(a))$ of a, and the maximum "depth" of a in P, namely the maximum number of input prefixes that guard an occurrence of a in P. Observe that, when P is well typed, the depth $[\![P]\!]_a$ is always finite. We write $<$ for the usual lexicographic order on triples returned by $\mathsf{measure}(a, \Gamma, P)$. Note that $<$ is well founded, *i.e.* there are no infinite descending chains of triples related by $<$, because the type system for lock freedom restricts priorities to natural numbers.

Lemma B.9. *Let* **(1)** $\Gamma \vdash_1 P$ *and* **(2)** Γ *balanced and* **(3)** *either* $\mathsf{in}(a, P)$ *or* $\mathsf{out}(a, P)$. *Then there exists* Q *such that* $P \rightarrow^* Q$ *and* $\mathsf{sync}(a, Q)$.

Proof. Without loss of generality we can make the following assumptions:

- that $\neg\mathsf{sync}(a, P)$, otherwise we conclude immediately by taking $Q = P$;
- that P has no restrictions at the top level. Indeed, if $P \equiv (\nu b)P'$, from **(1)** and [T-NEW] we deduce $\Gamma, b : \#^\iota[t]_m^n \vdash_1 P'$ where $\Gamma, b : \#^\iota[t]_m^n$ is balanced, so we could reason on P';
- that P has no $\texttt{let}$ or $\texttt{case}$ at the top level. If this is the case, then from **(1)** we can always reduce P to a well-typed process that contains no top level $\texttt{let}$'s or $\texttt{case}$'s.

Note that, from **(2)**, we know that the domain of Γ contains only channels (no variables). Therefore, any expression in any top-level output of P is evaluated.

We proceed by induction, showing that if the result holds for every triple (b, Γ', P') such that $\mathsf{measure}(b, \Gamma', P') < \mathsf{measure}(a, \Gamma, P)$, then it holds also for the triple (a, Γ, P). We begin by analyzing a few cases that are impossible given the typing rules of processes:

- $P \equiv \mathscr{C}[a?(x).R]$ and $a \in \mathsf{fn}(R)$. Then $a : \#[t]_m^n \in \Gamma$. According to [T-IN] we have $\Gamma', x : \$^n t \vdash_1 R$ and $n < |\Gamma'|$. Then $a \notin \mathsf{dom}(\Gamma')$, hence this case is impossible.
- $P \equiv \mathscr{C}[a!\langle e\rangle]$ and $a \in \mathsf{fn}(e)$ and $a : \#^\omega[t] \in \Gamma$ and $\neg*\mathsf{in}(a, P)$. Then from [T-OUT*] we deduce $\Gamma' \vdash e : \$_1^n t$ and $0 \leq |t|$. From $\neg*\mathsf{in}(a, P)$ we know that e must contain the endpoint a with input capability, for otherwise it would occur at the top level of P. Hence $|t| = \bot$, which contradicts $0 \leq |t|$. Hence this case is impossible.
- $P \equiv \mathscr{C}[a!\langle e\rangle]$ and $a \in \mathsf{fn}(e)$ and $a : \#[t]_m^n \in \Gamma$. Then from [T-OUT] we deduce $\Gamma' \vdash e : \$_1^n t$ and $0 < |t|$. We deduce

$$\begin{aligned}
n &< n + |t| && \text{from } 0 < |t| \\
&= |\$_1^n t| && \text{by definition of } |\cdot| \\
&\leq n && \text{from } a \in \mathsf{fn}(e)
\end{aligned}$$

which is absurd, so this case is also impossible.

Then we consider the cases in which there is a pending operation on a but no immediate synchronization is possible. In every case we are able to extend the reduction of P to a state where the measure of a is strictly smaller than $\mathsf{measure}(a, \Gamma, P)$. This is enough for establishing the result, given that the order $<$ on measures is well founded.

- $P \equiv \mathscr{C}[b!\langle v\rangle \mid *b?(x).R]$ and $a \in \mathsf{fn}(v)$. Then $P \rightarrow \mathscr{C}[R\{v/x\} \mid *b?(x).R] = P'$. By Lemma III.3 we deduce that $\Gamma' \vdash_1 P'$ for some Γ' such that $\mathsf{tickets}(\Gamma'(a)) < \mathsf{tickets}(\Gamma(a))$, therefore $\mathsf{measure}(a, \Gamma', P') < \mathsf{measure}(a, \Gamma, P)$. We conclude by induction hypothesis.
- $P \equiv \mathscr{C}[b!\langle v\rangle]$ where $a \in \mathsf{fn}(v)$ and b is a linear channel. From [T-OUT] we deduce $|\Gamma(b)| < |\Gamma(a)|$. By induction hypothesis we derive $P \rightarrow^* \mathscr{C}'[b!\langle v\rangle \mid b?(x).R] \rightarrow \mathscr{C}'[R\{v/x\}] = P'$. By Lemma III.3 there exists Γ' such that $\Gamma' \vdash_1 P'$ and $\mathsf{tickets}(\Gamma'(a)) < \mathsf{tickets}(\Gamma(a))$. We conclude by induction hypothesis.

14

- $P \equiv \mathscr{C}[b?(x).R]$ where $a \in \mathsf{fn}(R) \setminus \{b\}$. From [T-IN] we deduce $|\Gamma(b)| < |\Gamma(a)|$. By induction hypothesis we derive $P \stackrel{\tau}{\Longrightarrow} \mathscr{C}'[b!\langle \mathsf{v}\rangle \mid b?(x).R] \stackrel{b}{\longrightarrow} \mathscr{C}'[R\{\mathsf{v}/x\}] = P'$ for some v. By Lemma III.3 we deduce $\Gamma' \vdash_1 P'$ for some Γ' such that $\Gamma \stackrel{b}{\Longrightarrow} \Gamma'$ and $\mathsf{measure}(a, \Gamma', P') < \mathsf{measure}(a, \Gamma, P)$ because a is guarded by one less input prefix. We conclude by induction hypothesis. □

Theorem B.10 (Theorem III.4). *The following properties hold:*

1) *If $\vdash_{\mathsf{DF}} P$, then P is deadlock free.*
2) *If $\vdash_{\mathsf{LF}} P$, then P is lock free.*

Proof. The two properties are straightforward consequences of Lemma B.7 and Lemma B.9, respectively. □

<div align="center">

APPENDIX C

EXTENDED EXAMPLE: MULTIPARTY SESSIONS

</div>

In this section we show that a whole class of protocols whose lock freedom is guaranteed *by design* can be implemented by processes that are well-typed according to our type system for lock freedom. The challenging aspect of this exercise is that in general the processes will be recursive and interleaving possibly blocking actions pertaining several different linear channels configured in a cyclic network topology. We proceed as follows. First of all, we define a language for specifying protocols whose lock freedom is granted. We do so by means of a simple variant of global types [1], [12], [13], [16] describing sequences of *interaction events* between a fixed number of *participants*. Then, we extract from global types a characteristic implementation of the participants involved in the protocol, using the process language we have defined in Section II. Finally, we show how to construct a typing environment for the participants so that they are well typed according to the type discipline of Section III.

Extensions. For technical convenience, we work with a slightly more general syntax of types:

$$t ::= \cdots \mid \mu\alpha_k\{\alpha_i := t_i\}_{i=1..n} \mid \$_m^n\, t$$

Basically we generalize recursions so that the type $\mu\alpha_k\{\alpha_i := t_i\}_{i\in I}$ simultaneously binds *all* the type variables α_i for $i \in I$ in *each* t_i for $i \in I$. The elected type variable α_k with $k \in I$ stands for the whole term. It is known that this generalized recursion operator has the same expressive power as the one that binds exactly one variable [9], but it is more convenient to work with in this section. The notion of type contractiveness can be easily adjusted to this generalization. Equality between types works as before by considering the infinite unfolding of recursions, in particular, $\mu\alpha_k\{\alpha_i := t_i\}_{i\in I} = t_k\{\mu\alpha_i\{\alpha_i := t_i\}_{i\in I}/\alpha_i\}_{i\in I}$ where $t\{s_i/\alpha_i\}_{i\in I}$ denotes the simultaneous substitution of every free occurrence of α_i in t with s_i. We also consider $\$_m^n$ as a type constructor rather than a function, because we must be able to denote the shifting of *possibly open* types, whereas the shifting function (7) is defined for closed types only. We extend equality between types with the laws of (7). Because of contractiveness, every type can be rewritten into an equivalent one without $\$_m^n$ constructors.

Global types. We assume a finite, totally ordered set $\mathscr{R}$ of ==participant tags== p, q, ..., A, B, ... that we use to uniquely identify the participants of a protocol. We also need a set of ==labels== ℓ, ... for identifying events occurring in protocols. ==*Events*== ε, ... and ==global types== G, ... are defined below:

$$\varepsilon ::= \mathsf{p} \to \mathsf{q}@\ell \qquad G ::= \alpha \mid \varepsilon.G \mid \varepsilon.\{G, G\} \mid \mu\alpha.G$$

An event $\mathsf{p} \to \mathsf{q}@\ell$ represents the communication of a message from participant p (the sender) to participant q (the receiver). We omit the type of the message, because it is irrelevant for our purposes, and we will think of events merely as synchronizations. We decorate events with a label ℓ that uniquely identifies them, although we will see that the same label may occur several times in a given global type. The global type $\varepsilon.G$ describes a protocol in which ε *may* occur before all the other events in G (the assurance that it will *necessarily* occur before them depends on other properties of the global type that are not captured by their syntactic structure). The global type $\varepsilon.\{G_1, G_2\}$ describes a protocol in which the sender in ε communicates a binary decision to the receiver in ε. Depending on the outcome of the decision, the protocol evolves as either G_1 or G_2. Global types α and $\mu\alpha.G$ are used for building recursive protocols. We do not impose any contractiveness conditions on global types, which we treat as purely syntactic objects. In fact, we abbreviate $\mu\alpha.\alpha$ as end to denote the terminated protocol in which no event occurs. We write $\mathsf{fv}(G)$ for the set of free type variables in G, defined as expected.

Not every global type describes a sensible protocol. For example,

$$G_a \stackrel{\text{def}}{=} \mathsf{A} \to \mathsf{B}.\{\mathsf{C} \to \mathsf{D}.\mathsf{end}, \mathsf{D} \to \mathsf{C}.\mathsf{end}\}$$

specifies that either C sends a message to D or D sends a message to C depending on the decision taken by A. The problem is that neither C nor D are aware of such decision, for A communicates it only to B, so there is no way C and D can behave differently in the two branches of the choice. We therefore need to identify a class of global types that describe sensible protocols, and that we call *realizable*. An ==edge== $\mathsf{p}\wr\mathsf{q}$ is a set $\{\mathsf{p}, \mathsf{q}\}$ with $\mathsf{p} \neq \mathsf{q}$ representing a communication channel between two participants.

We let $\text{edge}(\text{p} \to \text{q}@\ell) = \text{p} \wr \text{q}$. A *sort* is a set of edges and tells us which pairs of participants are supposed to communicate. Since we will analyze possibly open global types, we define the sort of a global type relative to a sort environment that maps type variables to sorts:

Definition C.1. A *sort environment* Σ is a finite map from type variables to sorts. We say that Σ is a sort environment for G if $\text{fv}(G) \subseteq \text{dom}(\Sigma)$. The *sort* of G relative to a sort environment Σ for G is $\Sigma(G) \stackrel{\text{def}}{=} \{\text{edge}(\varepsilon) \mid \varepsilon \text{ occurs in } G\} \cup \bigcup_{\alpha \in \text{fv}(G)} \Sigma(\alpha)$.

We write $\emptyset$ for the empty sort environment and $\text{upd}(\Sigma, \alpha, G)$ for the updated sort environment $\Sigma, \alpha : \Sigma(\mu\alpha.G)$. We can now define the *projection* $\downarrow(\Sigma, G, \text{p} \wr \text{q})$ of a global type G with respect to a sort environment Σ and an edge $\text{p} \wr \text{q}$. Intuitively, $\downarrow(\Sigma, G, \text{p} \wr \text{q})$ extracts from G the sub-protocol that concerns only participants p and q. Formally:

$$
\begin{aligned}
\downarrow(\Sigma, \alpha, \text{p} \wr \text{q}) &\stackrel{\text{def}}{=} \alpha \\
\downarrow(\Sigma, \varepsilon.G, \text{p} \wr \text{q}) &\stackrel{\text{def}}{=} \varepsilon.\downarrow(\Sigma, G, \text{p} \wr \text{q}) && \text{if } \text{p} \wr \text{q} = \text{edge}(\varepsilon) \\
\downarrow(\Sigma, \varepsilon.G, \text{p} \wr \text{q}) &\stackrel{\text{def}}{=} \downarrow(\Sigma, G, \text{p} \wr \text{q}) && \text{if } \text{p} \wr \text{q} \neq \text{edge}(\varepsilon) \\
\downarrow(\Sigma, \varepsilon.\{G_1, G_2\}, \text{p} \wr \text{q}) &\stackrel{\text{def}}{=} \varepsilon. \begin{Bmatrix} \downarrow(\Sigma, G_1, \text{p} \wr \text{q}) \\ \downarrow(\Sigma, G_2, \text{p} \wr \text{q}) \end{Bmatrix} && \text{if } \text{p} \wr \text{q} = \text{edge}(\varepsilon) \\
\downarrow(\Sigma, \varepsilon.\{G_1, G_2\}, \text{p} \wr \text{q}) &\stackrel{\text{def}}{=} \downarrow(\Sigma, G_i, \text{p} \wr \text{q}) && \text{if } \text{p} \wr \text{q} \neq \text{edge}(\varepsilon) \\
\downarrow(\Sigma, \mu\alpha.G, \text{p} \wr \text{q}) &\stackrel{\text{def}}{=} \mu\alpha.\downarrow(\text{upd}(\Sigma, \alpha, G), G, \text{p} \wr \text{q})
\end{aligned}
$$

Note that projection is a *partial* function. It may be ill defined if the projections of two branches of a choice with respect to the same edge $\text{p} \wr \text{q}$ differ (fifth equation). For example, $\downarrow(\emptyset, G_a, \text{C} \wr \text{D})$ is ill defined because it can be either $\text{C} \to \text{D}.\text{end}$ or $\text{D} \to \text{C}.\text{end}$. Note the role played by labels in determining the existence of a projection. For example,

$$
G_b \stackrel{\text{def}}{=} \text{A} \to \text{B}@\ell_1. \begin{Bmatrix} \text{B} \to \text{A}@\ell_2.\text{C} \to \text{D}@\ell_3.\text{end} \\ \text{C} \to \text{D}@\ell_3.\text{B} \to \text{A}@\ell_4.\text{end} \end{Bmatrix}
$$

is projectable with respect to $\text{C} \wr \text{D}$ because the two occurrences of $\text{C} \to \text{D}@\ell_3$ in the two branches have the same label. They must in fact represent the same event, since C and D are unaware of the choice taken by A. On the contrary, $\text{B} \to \text{A}@\ell_2$ and $\text{B} \to \text{A}@\ell_4$ may have different labels because both A and B are aware of the choice taken by A.

Projectability alone is not a sufficient condition for realizability. We must also take into account the ordering of events induced by the structure of global types. For instance

$$
G_c \stackrel{\text{def}}{=} \text{A} \to \text{B}@\ell_1. \begin{Bmatrix} \text{E} \to \text{F}@\ell_2.\text{C} \to \text{D}@\ell_3.\text{end} \\ \text{C} \to \text{D}@\ell_3.\text{E} \to \text{F}@\ell_2.\text{end} \end{Bmatrix}
$$

is projectable but not realizable: the two occurrences of $\text{C} \to \text{D}$ and $\text{E} \to \text{F}$ must be tagged with the same label in order for $\downarrow(\emptyset, G_c, \text{C} \wr \text{D})$ and $\downarrow(\emptyset, G_c, \text{E} \wr \text{D})$ to be well defined, and yet G_c specifies that the order of these events depends on the decision taken by A, which none of C, D, E, and F is aware of. To capture these inconsistencies we reason on the relation $\prec_G$ induced by the structure of G such that $\ell \prec_G \ell'$ holds if and only if the event with label ℓ immediately precedes the event with label ℓ' in G. We denote by $\prec_G^+$ the transitive closure of $\prec_G$ and we require $\prec_G^+$ to be irreflexive for G to be realizable.

Definition C.2 (realizable global type). We say that G is *realizable* if **(1)** $\downarrow(\emptyset, G, \text{p} \wr \text{q})$ is well defined for every $\text{p} \wr \text{q}$ and **(2)** the relation $\prec_G^+$ is irreflexive.

According to this definition, G_b is realizable but G_c is not, because $\ell_2 \prec_{G_c}^+ \ell_2$. Note that the relation $\prec_G$ is solely determined by the syntactic structure of G regardless of recursions and type variables occurring in G. No non-trivial recursive global type would be realizable had we considered, for example, $\mu\alpha.\text{A} \to \text{B}@\ell.\alpha$ equivalent to its own unfolding $\text{A} \to \text{B}@\ell.\mu\alpha.\text{A} \to \text{B}@\ell.\alpha$.

Characteristic participants. We now show how to build a set of processes that implement any realizable global type. Every participant p in a global type G will be modeled as a sequential process $\mathbf{P}(\emptyset, G, \text{p})$ that uses a tuple $\mathbf{x}_{[\text{p}]} = x_{\text{pq}_1}, \ldots, x_{\text{pq}_n}$ of variables, where $\{\text{q}_1, \ldots, \text{q}_n\} = \mathscr{R} \setminus \{\text{p}\}$ and the q_i's appear according to the total order on $\mathscr{R}$ (the total order serves only so that $\mathbf{x}_{[\text{p}]}$ is a uniquely determined tuple). In particular, x_{pq} is the channel that connects p to q. $\mathbf{P}(\Sigma, G, \text{p})$ is defined inductively on G by the equations in the top half of Table 3. Despite the daunting look of the definition, the idea is simple: every event $\text{A} \to \text{B}$ to which p participates determines either an input or an output action on the channel x_{pq}. When p is the receiver, a message is read from x_{pq} and is given name x_{pq}, because the message is actually the continuation channel on which the conversation between p and q continues. When p is the sender, it sends to q the continuation for the subsequent communications on the edge $\text{p} \wr \text{q}$. This is done through the auxiliary function $\mathbf{C}(\Sigma, \text{p}, \text{q}, id, G)$ where id is the identity function: the continuation is an actual channel a if $\text{p} \wr \text{q}$ occurs in $\Sigma(G)$, that is if p and q will communicate again in the future, or the dummy number 0 if no other interaction between p and q is expected. When the event implies a choice $\varepsilon.\{G_1, G_2\}$, we make the (arbitrary) decision that the sender always chooses the left branch G_1. We use injection to encode the decision in the sender and `case` to decode the decision in the receiver and we use the same auxiliary function $\mathbf{C}$ for sending the decision, to which we pass the injector `inl` to wrap the continuation. Recursions $\mu\alpha.G$ are implemented by associating the type variable α with

TABLE 3: Global type projections.

$$
\begin{aligned}
\mathbf{P}(\Sigma, \alpha, \mathtt{p}) &\stackrel{\text{def}}{=} c_\alpha!\langle \mathbf{x}_{[\mathtt{p}]}\rangle \\
\mathbf{P}(\Sigma, \mathtt{A}\to\mathtt{B}.G, \mathtt{A}) &\stackrel{\text{def}}{=} \mathbf{C}(\Sigma, \mathtt{A}, \mathtt{B}, id, G) \\
\mathbf{P}(\Sigma, \mathtt{A}\to\mathtt{B}.G, \mathtt{B}) &\stackrel{\text{def}}{=} x_{\mathtt{BA}}?(x_{\mathtt{BA}}).\mathbf{P}(\Sigma, G, \mathtt{B}) \\
\mathbf{P}(\Sigma, \mathtt{A}\to\mathtt{B}.G, \mathtt{p}) &\stackrel{\text{def}}{=} \mathbf{P}(\Sigma, G, \mathtt{p}) && \mathtt{p}\notin \mathtt{A}\wr\mathtt{B} \\
\mathbf{P}(\Sigma, \mathtt{A}\to\mathtt{B}.\{G_1, G_2\}, \mathtt{A}) &\stackrel{\text{def}}{=} \mathbf{C}(\Sigma, \mathtt{A}, \mathtt{B}, \mathtt{inl}, G_1) \\
\mathbf{P}(\Sigma, \mathtt{A}\to\mathtt{B}.\{G_1, G_2\}, \mathtt{B}) &\stackrel{\text{def}}{=} x_{\mathtt{BA}}?(y).\mathtt{case}\ y\ \{\mathtt{inl}\ x_{\mathtt{BA}}\Rightarrow \mathbf{P}(\Sigma, G_1, \mathtt{B}), \mathtt{inr}\ x_{\mathtt{BA}}\Rightarrow \mathbf{P}(\Sigma, G_2, \mathtt{B})\} \\
\mathbf{P}(\Sigma, \mathtt{A}\to\mathtt{B}.\{G_1, G_2\}, \mathtt{p}) &\stackrel{\text{def}}{=} \mathbf{P}(\Sigma, G_i, \mathtt{p}) && \mathtt{p}\notin \mathtt{A}\wr\mathtt{B} \\
\mathbf{P}(\Sigma, \mu\alpha.G, \mathtt{p}) &\stackrel{\text{def}}{=} (\nu c_\alpha)(c_\alpha!\langle \mathbf{x}_{[\mathtt{p}]}\rangle \mid *c_\alpha?(\mathbf{x}_{[\mathtt{p}]}).\mathbf{P}(\mathsf{upd}(\Sigma, \alpha, G), G, \mathtt{p})) \\
\mathbf{C}(\Sigma, \mathtt{p}, \mathtt{q}, f, G) &\stackrel{\text{def}}{=} (\nu a)(x_{\mathtt{pq}}!\langle f\ a\rangle \mid \mathtt{let}\ x_{\mathtt{pq}}=a\ \mathtt{in}\ \mathbf{P}(\Sigma, G, \mathtt{p})) && \mathtt{p}\wr\mathtt{q}\in\Sigma(G) \\
\mathbf{C}(\Sigma, \mathtt{p}, \mathtt{q}, f, G) &\stackrel{\text{def}}{=} x_{\mathtt{pq}}!\langle f\ 0\rangle \mid \mathtt{let}\ x_{\mathtt{pq}}=0\ \mathtt{in}\ \mathbf{P}(\Sigma, G, \mathtt{p}) && \mathtt{p}\wr\mathtt{q}\notin\Sigma(G)
\end{aligned}
$$

$$
\begin{aligned}
\mathbf{T}(\Sigma, \alpha, \mathtt{p}, \mathtt{q}, n) &\stackrel{\text{def}}{=} \begin{cases} \$_1^0\, \alpha_{\mathtt{pq}} & \mathtt{p}\wr\mathtt{q}\in\Sigma(\alpha) \\ \mathtt{int} & \text{otherwise} \end{cases} \\[2mm]
\mathbf{T}(\Sigma, \mathtt{A}\to\mathtt{B}@\ell.G, \mathtt{p}, \mathtt{q}, n) &\stackrel{\text{def}}{=} \begin{cases} \$^{\ell-n}\, ![\mathbf{T}(\Sigma, G, \mathtt{q}, \mathtt{p}, \ell)] & \mathtt{p}, \mathtt{q}=\mathtt{A}, \mathtt{B} \\ \$^{\ell-n}\, ?[\mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, \ell)] & \mathtt{p}, \mathtt{q}=\mathtt{B}, \mathtt{A} \\ \$^{\ell-n}\, \mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, \ell) & \text{otherwise} \end{cases} \\[2mm]
\mathbf{T}(\Sigma, \mathtt{A}\to\mathtt{B}@\ell.\{G_1, G_2\}, \mathtt{p}, \mathtt{q}, n) &\stackrel{\text{def}}{=} \begin{cases} \$^{\ell-n}\, ![\mathbf{T}(\Sigma, G_1, \mathtt{q}, \mathtt{p}, \ell)\oplus \mathbf{T}(\Sigma, G_2, \mathtt{q}, \mathtt{p}, \ell)] & \mathtt{p}, \mathtt{q}=\mathtt{A}, \mathtt{B} \\ \$^{\ell-n}\, ?[\mathbf{T}(\Sigma, G_1, \mathtt{p}, \mathtt{q}, \ell)\oplus \mathbf{T}(\Sigma, G_2, \mathtt{p}, \mathtt{q}, \ell)] & \mathtt{p}, \mathtt{q}=\mathtt{B}, \mathtt{A} \\ \$^{\ell-n}\, \mathbf{T}(\Sigma, G_i, \mathtt{p}, \mathtt{q}, \ell) & \text{otherwise} \end{cases} \\[2mm]
\mathbf{T}(\Sigma, \mu\alpha.G, \mathtt{p}, \mathtt{q}, n) &\stackrel{\text{def}}{=} \$_1^0\, \mu\alpha_{\mathtt{pq}} \begin{Bmatrix} \alpha_{\mathtt{pq}}=\mathbf{T}(\mathsf{upd}(\Sigma, \alpha, G), G, \mathtt{p}, \mathtt{q}, n) \\ \alpha_{\mathtt{qp}}=\mathbf{T}(\mathsf{upd}(\Sigma, \alpha, G), G, \mathtt{q}, \mathtt{p}, n) \end{Bmatrix}
\end{aligned}
$$

an unlimited channel c_α: the replicated input on c_α corresponds to defining a recursive process; the output on c_α corresponds to invoking the recursive process. The whole tuple $\mathbf{x}_{[\mathtt{p}]}$ is transmitted at each invocation.

For instance, given $G_d \stackrel{\text{def}}{=} \mu\alpha.\mathtt{A}\to\mathtt{B}.\mathtt{B}\to\mathtt{C}.\mathtt{C}\to\mathtt{A}.\alpha$ and assuming $\mathscr{R}=\{\mathtt{A}, \mathtt{B}, \mathtt{C}\}$, we have

$$
\begin{aligned}
\mathbf{P}(\emptyset, G_d, \mathtt{A}) = \\
(\nu c_\alpha)(c_\alpha!\langle x_{\mathtt{AB}}, x_{\mathtt{AC}}\rangle \mid \\
*c_\alpha?(x_{\mathtt{AB}}, x_{\mathtt{AC}}).(\nu a)(x_{\mathtt{AB}}!\langle a\rangle \mid \mathtt{let}\ x_{\mathtt{AB}}=a\ \mathtt{in} \\
x_{\mathtt{AC}}?(x_{\mathtt{AC}}).c_\alpha!\langle x_{\mathtt{AB}}, x_{\mathtt{AC}}\rangle))
\end{aligned}
$$

Typing participants. We now arrive at the main point of this section, namely the definition of a suitable type environment $\Gamma_{\mathtt{p}} \stackrel{\text{def}}{=} \{x_{\mathtt{pq}} : t_{\mathtt{pq}}\}_{\mathtt{q}\in\mathscr{R}\setminus\{\mathtt{p}\}}$ such that the process $\mathbf{P}(\emptyset, G, \mathtt{p})$ is well typed in $\Gamma_{\mathtt{p}}$. Determining the input/output behavior of $\mathbf{P}(\emptyset, G, \mathtt{p})$ with respect to a channel $x_{\mathtt{pq}}$ is easy: this information is written in the global type. But for $\mathbf{P}(\emptyset, G, \mathtt{p})$ to be well typed, we must also make sure that blocking actions respect the priority of channels (see rule [T-IN]) and that delegations are allowed by the tickets of channels (see [T-OUT] and [T-OUT*]). We determine priorities using the event order $\prec_G$ and tickets looking at the structure of $\mathbf{P}(\emptyset, G, \mathtt{p})$. We know that every realizable global type G has an irreflexive event order $\prec_G^+$, meaning that the events in G form a DAG. Therefore, there exists a topological ordering ord from labels in G to positive natural numbers such that $\ell \prec_G \ell'$ implies $0 < \mathsf{ord}(\ell) < \mathsf{ord}(\ell')$. In a sense, $\mathsf{ord}(\ell)$ is the abstract time at which the event labeled ℓ occurs. The idea is to use $\mathsf{ord}(\ell)$ as the priority for the channel types when projecting the action labeled by ℓ. In what follows, we simplify the notation writing just ℓ in place of $\mathsf{ord}(\ell)$.

Now, the bottom half of Table 3 defines $\mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, n)$ as the type of $x_{\mathtt{pq}}$ in the global type G relative to an arbitrary offset n that is smaller than any label occurring in G. Given how $\mathbf{P}(\Sigma, G, \mathtt{p})$ uses the channel $x_{\mathtt{pq}}$, the definition of $\mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, n)$ is hopefully self-explanatory except for a few twists that we comment on here: The type $\mathbf{T}(\Sigma, \alpha, \mathtt{p}, \mathtt{q}, n)$ is the type variable $\alpha_{\mathtt{pq}}$ only if the edge $\mathtt{p}\wr\mathtt{q}$ is used in the recursion marked by the type variable α; otherwise, it is $\mathtt{int}$ to denote a dummy natural number. Priorities are determined as the difference between the label ℓ of the topmost action in G and the current offset n; after the action, the offset becomes ℓ. We have $\mathbf{T}(\Sigma, \mathtt{A}\to\mathtt{B}@\ell.G, \mathtt{A}, \mathtt{B}, n) = \$^{\ell-n}\, ![\mathbf{T}(\Sigma, G, \mathtt{B}, \mathtt{A}, \ell)]$ where $\mathtt{A}$ and $\mathtt{B}$ are *swapped* within $![\cdot]$. The swapping is motivated by the fact that $\mathtt{A}$ sends to $\mathtt{B}$ a channel whose type determines how the channel is used by $\mathtt{B}$. The same phenomenon occurs in determining the type of a choice. Then, because of the definition of $\mathbf{T}(\Sigma, \alpha, \mathtt{p}, \mathtt{q}, n)$, it may happen that the type $\mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, n)$ contains both the $\alpha_{\mathtt{pq}}$ and the $\alpha_{\mathtt{qp}}$ type variables. For this reason, $\mathbf{T}(\Sigma, \mu\alpha.G, \mathtt{p}, \mathtt{q}, n)$ binds both $\alpha_{\mathtt{pq}}$ and $\alpha_{\mathtt{qp}}$, taking advantage of the generalized recursion for channel types. We apply $\$_1^0$ in the two places where a channel is delegated, in accordance with the definition of $\mathbf{P}(\Sigma, G, \mathtt{p})$. Finally, note that $\mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, n)$ is well defined whenever G is realizable.

For example, considering again G_d above, we have

$$\mathbf{T}(\emptyset, G_d, \mathtt{A}, \mathtt{B}, 0) = \$_1^0 \, \mu\alpha_{\mathtt{AB}}. \begin{cases} \alpha_{\mathtt{AB}} := \, ! \, [\$_1^{\ell_3 - \ell_1} \, \alpha_{\mathtt{BA}}]^{\ell_1} 0 \\ \alpha_{\mathtt{BA}} := \, ? \, [\$_1^{\ell_3 - \ell_1} \, \alpha_{\mathtt{BA}}]^{\ell_1} 0 \end{cases}$$

assuming that ℓ_1, ℓ_2, and ℓ_3 label the three events in G_d.

The main result of this section states that the $\mathbf{P}(\emptyset, G, \mathtt{p})$'s are well typed and so is their parallel composition:

Theorem C.3. *Let G be realizable. Then for every $\mathtt{p} \in \mathscr{R}$ we have $\{x_{\mathtt{pq}} : \mathbf{T}(\emptyset, G, \mathtt{p}, \mathtt{q}, 0) \mid \mathtt{q} \in \mathscr{R} \setminus \{\mathtt{p}\}\} \vdash_1 \mathbf{P}(\emptyset, G, \mathtt{p})$.*

Note that a global type with two participants is always realizable. This corresponds to a binary session.

APPENDIX D
PROOFS FOR APPENDIX C

We write $\bar{p}$ for $\{?, !\} \setminus p$. We write $n < G$ if $n < \mathsf{ord}(\ell)$ for every ℓ occurring in G.

Lemma D.1. *Let $n < G$ and $\mathbf{T}(\Sigma, G, \mathtt{p}, \mathtt{q}, n) = t$. Then:*
1) $\mathtt{p} \wr \mathtt{q} \notin \Sigma(G)$ *implies* $t = \mathtt{int}$;
2) $\mathtt{p} \wr \mathtt{q} \notin G$ *and* $\mathtt{p} \wr \mathtt{q} \in \Sigma(\alpha)$ *and* $\alpha \in \mathsf{fv}(G)$ *implies* $t = \$_k^h \, \alpha_{\mathtt{pq}}$;
3) $\mathtt{p} \wr \mathtt{q} \in G$ *implies* $t = p[s]_k^h$ *and* $\mathbf{T}(\Sigma, G, \mathtt{q}, \mathtt{p}, n) = \bar{p}[s]_k^h$ *for some p and $h > 0$.*

Proof. By induction on G.

- ($G = \alpha$) We have:
  1) ($\mathtt{p} \wr \mathtt{q} \notin \Sigma(G)$) Then $t = \mathtt{int}$ by definition of $\mathbf{T}$.
  2) ($\mathtt{p} \wr \mathtt{q} \in \Sigma(\alpha)$) Then $t = \$_1^0 \, \alpha_{\mathtt{pq}}$ by definition of $\mathbf{T}$ and we conclude by taking $h = 0$ and $k = 1$.
  3) ($\mathtt{p} \wr \mathtt{q} \in G$) This case is impossible.
- ($G = \mathtt{A} \to \mathtt{B} @ \ell.G'$) We have:
  1) ($\mathtt{p} \wr \mathtt{q} \notin \Sigma(G)$) Then $\mathtt{p} \wr \mathtt{q} \neq \mathtt{A} \wr \mathtt{B}$ and $\mathtt{p} \wr \mathtt{q} \notin \Sigma(G')$ and $t = \$^{\ell - n} \, \mathbf{T}(\Sigma, G', \mathtt{p}, \mathtt{q}, \ell) = \$^{\ell - n} \, \mathtt{int} = \mathtt{int}$ by definition of $\mathbf{T}$, by induction hypothesis, and by equality on types.
  2) ($\mathtt{p} \wr \mathtt{q} \notin G$ and $\alpha \in \mathsf{fv}(G)$ and $\mathtt{p} \wr \mathtt{q} \in \Sigma(G)$) Then $\mathtt{p} \wr \mathtt{q} \neq \mathtt{A} \wr \mathtt{B}$ and $\mathtt{p} \wr \mathtt{q} \notin G'$ and $\alpha \in \mathsf{fv}(G')$. We have $t = \$^{\ell - n} \, \mathbf{T}(\Sigma, G', \mathtt{p}, \mathtt{q}, \ell) = \$^{\ell - n} \$_m^m \, \alpha_{\mathtt{pq}} = \$_k^{\ell - n + m} \, \alpha_{\mathtt{pq}}$ by definition of $\mathbf{T}$ and by induction hypothesis. We conclude by taking $h = \ell - n + m$.
  3) ($\mathtt{p} \wr \mathtt{q} \in G$) Then either $\mathtt{p} \wr \mathtt{q} = \mathtt{A} \wr \mathtt{B}$ or $\mathtt{p} \wr \mathtt{q} \in G'$. If $\mathtt{p}, \mathtt{q} = \mathtt{A}, \mathtt{B}$, then $t = ![s]_k^h$ where $h = \ell - n$ and $k = 0$ and $s = \mathbf{T}(\Sigma, G', \mathtt{q}, \mathtt{p}, \ell)$ and we conclude by observing that $h > 0$ from the hypothesis $n < G$. If $\mathtt{p}, \mathtt{q} = \mathtt{B}, \mathtt{A}$ we conclude similarly. If $\mathtt{p} \wr \mathtt{q} \neq \mathtt{A} \wr \mathtt{B}$, then by induction hypothesis we have $\mathbf{T}(\Sigma, G', \mathtt{p}, \mathtt{q}, \ell) = p[s]_k^m$ for some $m > 0$ and we conclude $t = p[s]_k^h$ by taking $h = \ell - n + m$.
- ($G = \mathtt{A} \to \mathtt{B} @ \ell.\{G_1, G_2\}$) This case is similar to the previous one, the only interesting sub-case is when $\mathtt{p} \wr \mathtt{q} \notin G$ and $\alpha \in \mathsf{fv}(G)$ and $\mathtt{p} \wr \mathtt{q} \in \Sigma(G)$. By definition of global type projection we have $t = \mathbf{T}(\Sigma, G_i, \mathtt{p}, \mathtt{q}, \ell)$ for every $i = 1, 2$, therefore we deduce $\alpha \in \mathsf{fv}(G_1) \cap \mathsf{fv}(G_2)$. We conclude by induction hypothesis as in the previous case.
- ($G = \mu\alpha.G'$) Let

$$\begin{aligned} \Sigma' &\stackrel{\text{def}}{=} \mathsf{upd}(\Sigma, \alpha, G') \\ t_{\mathtt{pq}} &\stackrel{\text{def}}{=} \mu\alpha_{\mathtt{pq}}\{\alpha_{\mathtt{pq}} := \mathbf{T}(\Sigma', G', \mathtt{p}, \mathtt{q}, n), \alpha_{\mathtt{qp}} := \mathbf{T}(\Sigma', G', \mathtt{q}, \mathtt{p}, n)\} \\ t_{\mathtt{qp}} &\stackrel{\text{def}}{=} \mu\alpha_{\mathtt{qp}}\{\alpha_{\mathtt{pq}} := \mathbf{T}(\Sigma', G', \mathtt{p}, \mathtt{q}, n), \alpha_{\mathtt{qp}} := \mathbf{T}(\Sigma', G', \mathtt{q}, \mathtt{p}, n)\} \\ \sigma &\stackrel{\text{def}}{=} \{t_{\mathtt{pq}}, t_{\mathtt{qp}}/\alpha_{\mathtt{pq}}, \alpha_{\mathtt{qp}}\} \end{aligned}$$

  and observe that Σ' is a sort environment for G' and $t = \$_1^0 \, t_{\mathtt{pq}} = \$_1^0 \, \mathbf{T}(\Sigma', G', \mathtt{p}, \mathtt{q}, n)\sigma$ and $\mathbf{T}(\Sigma, G, \mathtt{q}, \mathtt{p}, n) = \$_1^0 \, t_{\mathtt{qp}} = \$_1^0 \, \mathbf{T}(\Sigma', G', \mathtt{q}, \mathtt{p}, n)\sigma$. We distinguish three sub-cases:
  1) ($\mathtt{p} \wr \mathtt{q} \notin \Sigma(G)$) Then $\mathtt{p} \wr \mathtt{q} \notin \Sigma'(G')$ and we conclude $t = \mathtt{int}$.
  2) ($\mathtt{p} \wr \mathtt{q} \notin G$ and $\mathtt{p} \wr \mathtt{q} \in \Sigma(\beta)$ and $\beta \in \mathsf{fv}(G)$) Then $\mathtt{p} \wr \mathtt{q} \notin G'$ and $\beta \neq \alpha$ and $\beta \in \mathsf{fv}(G')$. By induction hypothesis we deduce $\mathbf{T}(\Sigma', G', \mathtt{p}, \mathtt{q}, n) = \$_m^h \, \beta_{\mathtt{pq}}$. We conclude by taking $k = m + 1$ and observing that $t = \$_k^h \, \beta_{\mathtt{pq}}$.
  3) ($\mathtt{p} \wr \mathtt{q} \in G$) Then $\mathtt{p} \wr \mathtt{q} \in G'$ and by induction hypothesis we deduce $\mathbf{T}(\Sigma', G', \mathtt{p}, \mathtt{q}, n) = p[s]_m^h$ and $\mathbf{T}(\Sigma', G', \mathtt{q}, \mathtt{p}, n) = \bar{p}[s]_m^h$ for some $h > 0$. We conclude by taking $k = m + 1$ and observing that $t = p[s\sigma]_k^h$ and $\mathbf{T}(\Sigma, G, \mathtt{q}, \mathtt{p}, n) = \bar{p}[s\sigma]_k^h$. $\qquad \square$

Definition D.2 (Σ-substitution). A Σ-substitution is a finite map σ such that:
1) $\mathsf{dom}(\sigma) = \{\alpha_{\mathtt{pq}} \mid \alpha \in \mathsf{dom}(\Sigma) \wedge \mathtt{p} \wr \mathtt{q} \subseteq \mathscr{R}\}$;
2) $\sigma(\alpha_{\mathtt{pq}}) = \sigma(\alpha_{\mathtt{qp}}) = \mathtt{int}$ for every $\mathtt{p} \wr \mathtt{q} \notin \Sigma(\alpha)$;
3) $\sigma(\alpha_{\mathtt{pq}}) = p[t]_n^m$ and $\sigma(\alpha_{\mathtt{qp}}) = \bar{p}[t]_n^m$ and $m > 0$ for every $\mathtt{p} \wr \mathtt{q} \in \Sigma(\alpha)$.

18

Lemma D.3. *If $\Sigma' = \Sigma, \alpha : \Sigma(\mu\alpha.G)$, then $\Sigma(\mu\alpha.G) = \Sigma'(G)$.*

Proof. We have

$$
\begin{aligned}
\Sigma(\mu\alpha.G) &= \Sigma(\mu\alpha.G) \underbrace{\cup\, \Sigma(\mu\alpha.G)}_{\alpha \in \mathsf{fv}(G)} && \text{property of set union} \\
&= \mathsf{sort}(\mu\alpha.G) \cup \bigcup_{\beta \in \mathsf{fv}(\mu\alpha.G)} \Sigma(\beta) \underbrace{\cup\, \Sigma(\mu\alpha.G)}_{\alpha \in \mathsf{fv}(G)} && \text{definition of } \Sigma(\mu\alpha.G) \\
&= \mathsf{sort}(\mu\alpha.G) \cup \bigcup_{\beta \in \mathsf{fv}(G)} \Sigma'(\beta) && \mathsf{fv}(\mu\alpha.G) = \mathsf{fv}(G) \setminus \{\alpha\} \\
&= \mathsf{sort}(G) \cup \bigcup_{\beta \in \mathsf{fv}(G)} \Sigma'(\beta) && \mathsf{sort}(\mu\alpha.G) = \mathsf{sort}(G) \\
&= \Sigma'(G) && \text{definition of } \Sigma'(G)
\end{aligned}
$$

$\square$

Lemma D.4. *Let*
- *σ be a Σ-substitution;*
- *$n < G$;*
- *$\Sigma' = \Sigma, \alpha : \Sigma(\mu\alpha.G)$;*
- *$s_{\mathsf{pq}} = \mu\alpha_{\mathsf{pq}}.\{\alpha_{\mathsf{pq}} := \mathbf{T}(\Sigma', G, \mathsf{p}, \mathsf{q}, n), \alpha_{\mathsf{qp}} := \mathbf{T}(\Sigma', G, \mathsf{q}, \mathsf{p}, n)\}$;*
- *$\sigma' = \sigma, \{s_{\mathsf{pq}}\sigma/\alpha_{\mathsf{pq}}\}_{\mathsf{p}\wr\mathsf{q} \subseteq \mathscr{R}}$.*

Then σ' is a Σ'-substitution.

Proof. We must prove the three conditions of Definition D.2. The first condition is satisfied because $\mathsf{dom}(\sigma') = \mathsf{dom}(\sigma) \cup \{\alpha_{\mathsf{pq}} \mid \mathsf{p}\wr\mathsf{q} \subseteq \mathscr{R}\}$, so let us consider conditions 2 and 3. For every type variable $\beta \neq \alpha$ these conditions are trivially satisfied from the hypothesis that σ is a Σ-substitution, therefore we focus on α. Observe that, by Lemma D.3, we have $\Sigma'(\alpha) = \Sigma(\mu\alpha.G) = \Sigma'(G)$.

Regarding condition 2 of Definition D.2, assume $\mathsf{p}\wr\mathsf{q} \notin \Sigma'(\alpha) = \Sigma'(G)$. We conclude

$$
\begin{aligned}
\sigma'(\alpha_{\mathsf{pq}}) &= s_{\mathsf{pq}}\sigma && \text{by definition of } \sigma' \\
&= \mathbf{T}(\Sigma', G, \mathsf{p}, \mathsf{q}, n)\{s_{\mathsf{pq}}, s_{\mathsf{qp}}/\alpha_{\mathsf{pq}}, \alpha_{\mathsf{qp}}\}\sigma && \text{by unfolding of } s_{\mathsf{pq}} \\
&= \mathtt{int}\{s_{\mathsf{pq}}, s_{\mathsf{qp}}/\alpha_{\mathsf{pq}}, \alpha_{\mathsf{qp}}\}\sigma && \text{by Lemma D.1} \\
&= \mathtt{int} && \text{by definition of type substitution}
\end{aligned}
$$

Regarding condition 3 of Definition D.2, assume $\mathsf{p}\wr\mathsf{q} \in \Sigma'(\alpha) = \Sigma'(G)$. We distinguish two sub-cases:
- ($\mathsf{p}\wr\mathsf{q} \notin G$ and $\mathsf{p}\wr\mathsf{q} \in \Sigma'(\beta)$ and $\beta \in \mathsf{fv}(G)$) From $\mathsf{p}\wr\mathsf{q} \notin G$ we deduce $\beta \neq \alpha$, so

$$
\begin{aligned}
\sigma'(\alpha_{\mathsf{pq}}) &= \mathbf{T}(\Sigma', G, \mathsf{p}, \mathsf{q}, n)\{s_{\mathsf{pq}}, s_{\mathsf{qp}}/\alpha_{\mathsf{pq}}, \alpha_{\mathsf{qp}}\}\sigma && \text{by definition of } \sigma' \\
&= \$_k^h\, \beta_{\mathsf{pq}}\{s_{\mathsf{pq}}, s_{\mathsf{qp}}/\alpha_{\mathsf{pq}}, \alpha_{\mathsf{qp}}\}\sigma && \text{by Lemma D.1(2)} \\
&= \$_k^h\, \sigma(\beta_{\mathsf{pq}}) && \text{because } \beta \neq \alpha \\
&= \$_k^h\, p[s]_{k'}^{h'} && \text{because } \sigma \text{ is a } \Sigma\text{-substitution} \\
&= p[s]_{k+k'}^{h+h'} && \text{equality on types}
\end{aligned}
$$

  and similarly we obtain $\sigma'(\alpha_{\mathsf{pq}}) = \bar{p}[s]_{k+k'}^{h+h'}$.
- ($\mathsf{p}\wr\mathsf{q} \in G$) We have

$$
\begin{aligned}
\sigma'(\alpha_{\mathsf{pq}}) &= \mathbf{T}(\Sigma', G, \mathsf{p}, \mathsf{q}, n)\{s_{\mathsf{pq}}, s_{\mathsf{qp}}/\alpha_{\mathsf{pq}}, \alpha_{\mathsf{qp}}\}\sigma && \text{by definition of } \sigma' \\
&= p[t]_k^h && \text{by Lemma D.1(3)}
\end{aligned}
$$

  and similarly we obtain $\sigma'(\alpha_{\mathsf{qp}}) = \bar{p}[t]_k^h$. $\square$

Lemma D.5. *Let $n < G$ and σ be a Σ-substitution. Then:*
- *$\mathsf{p}\wr\mathsf{q} \notin \Sigma(G)$ implies $\mathbf{T}(\Sigma, G, \mathsf{p}, \mathsf{q}, n)\sigma = \mathtt{int}$;*
- *$\mathsf{p}\wr\mathsf{q} \in \Sigma(G)$ implies $\mathbf{T}(\Sigma, G, \mathsf{p}, \mathsf{q}, n)\sigma = p[t]_k^h$ and $\mathbf{T}(\Sigma, G, \mathsf{q}, \mathsf{p}, n)\sigma = \bar{p}[t]_k^h$ and $h > 0$.*

Proof. If $\mathsf{p}\wr\mathsf{q} \notin \Sigma(G)$, then by Lemma D.1 we deduce $\mathbf{T}(\Sigma, G, \mathsf{p}, \mathsf{q}, n) = \mathtt{int}$ and we conclude immediately. If $\mathsf{p}\wr\mathsf{q} \in \Sigma(G)$, then we distinguish two sub-cases:
- ($\mathsf{p}\wr\mathsf{q} \in G$) By Lemma D.1 we have $\mathbf{T}(\Sigma, G, \mathsf{p}, \mathsf{q}, n) = p[s]_k^h$ and $\mathbf{T}(\Sigma, G, \mathsf{q}, \mathsf{p}, n) = \bar{p}[s]_k^h$ and $h > 0$. We conclude by taking $t = s\sigma$.
- ($\mathsf{p}\wr\mathsf{q} \notin G$ and $\alpha \in \mathsf{fv}(G)$ and $\mathsf{p}\wr\mathsf{q} \in \Sigma(\alpha)$) By Lemma D.1 we have $\mathbf{T}(\Sigma, G, \mathsf{p}, \mathsf{q}, n) = \$_{k'}^{h'} \alpha_{\mathsf{pq}}$ and $\mathbf{T}(\Sigma, G, \mathsf{q}, \mathsf{p}, n) = \$_{k'}^{h'} \alpha_{\mathsf{qp}}$. By definition of Σ-substitution we deduce $\sigma(\alpha_{\mathsf{pq}}) = p[s]_{k''}^{h''}$ and $\sigma(\alpha_{\mathsf{qp}}) = \bar{p}[s]_{k''}^{h''}$ and $h'' > 0$. We conclude by taking $h = h' + h''$ and $k = k' + k''$. $\square$

19

We introduce the following abbreviations for referring to the environments used in the typing derivations of the following lemma:

$$\Theta_{[\Sigma,\mathrm{p}]} \;\stackrel{\text{def}}{=}\; \{c_\alpha : !^\omega[\alpha_{[\mathrm{p}]}] \mid \alpha \in \mathsf{dom}(\Sigma)\}$$
$$\Gamma_{[\Sigma,G,\mathrm{p},n]} \;\stackrel{\text{def}}{=}\; \{x_{\mathrm{pq}} : \mathbf{T}(\Sigma,G,\mathrm{p},\mathrm{q},n) \mid \mathrm{q} \in \mathscr{R}\setminus\{\mathrm{p}\}\} = \mathbf{x}_{[\mathrm{p}]} : \mathbf{T}(\Sigma,G,[\mathrm{p}],n)$$

Lemma D.6. *Let σ be a Σ-substitution and $n < G$. Then $\Theta_{[\Sigma,\mathrm{p}]}\sigma, \Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash_1 \mathbf{P}(\Sigma,G,\mathrm{p})$.*

Proof. By induction on G. We omit the cases for choices, since they are analogous to the ones for interactions.

$\boxed{G = \alpha}$ Let $s_{\mathrm{pq}} \stackrel{\text{def}}{=} \sigma(\alpha_{\mathrm{pq}})$. We have:

$$
\begin{aligned}
\mathbf{P}(\Sigma,G,\mathrm{p}) &= c_\alpha!\langle \mathbf{x}_{[\mathrm{p}]}\rangle & & \\
\Theta_{[\Sigma,\mathrm{p}]}\sigma &\ni c_\alpha : !^\omega[\alpha_{[\mathrm{p}]}]\sigma &=& c_\alpha : !^\omega[s_{[\mathrm{p}]}] \\
\Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma &= \mathbf{x}_{[\mathrm{p}]} : \mathbf{T}(\Sigma,G,[\mathrm{p}],n)\sigma &=& \mathbf{x}_{[\mathrm{p}]} : \$_1^0\, s_{[\mathrm{p}]}
\end{aligned}
$$

We derive:

$$
\frac{\Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash \mathbf{x}_{[\mathrm{p}]} : \$_1^0\, s_{[\mathrm{p}]}}{\Theta_{[\Sigma,\mathrm{p}]}\sigma, \Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash_1 c_\alpha!\langle \mathbf{x}_{[\mathrm{p}]}\rangle}\;[\text{T-OUT*}]
$$

where we use the fact that $0 \le |(|s_{\mathrm{pq}})$ which follows from the hypothesis that σ is a Σ-substitution.

$\boxed{G = \mu\alpha.G'}$ Let

$$
\begin{aligned}
\Sigma' &\stackrel{\text{def}}{=} \mathsf{upd}(\Sigma,\alpha,G') = \Sigma, \alpha : \Sigma(G) \\
s_{[\mathrm{p}]} &\stackrel{\text{def}}{=} s_{\mathrm{pq}_1} \times \cdots \times s_{\mathrm{pq}_n} \\
s_{\mathrm{pq}} &\stackrel{\text{def}}{=} \mu\alpha_{\mathrm{pq}}\{\alpha_{\mathrm{pq}} := \mathbf{T}(\Sigma',G',\mathrm{p},\mathrm{q},n), \alpha_{\mathrm{qp}} := \mathbf{T}(\Sigma',G',\mathrm{q},\mathrm{p},n)\} \\
\sigma' &\stackrel{\text{def}}{=} \sigma \circ \{s_{[\mathrm{p}]}/\alpha_{[\mathrm{p}]}\}
\end{aligned}
$$

and observe that σ' is a Σ'-substitution by Lemma D.4. We have:

$$
\begin{aligned}
\mathbf{P}(\Sigma,G,\mathrm{p}) &= (\nu c_\alpha)(c_\alpha!\langle \mathbf{x}_{[\mathrm{p}]}\rangle \mid *c_\alpha?(\mathbf{x}_{[\mathrm{p}]}).\mathbf{P}(\Sigma',G',\mathrm{p})) & & \\
\Theta_{[\Sigma',\mathrm{p}]}\sigma' &= (\Theta_{[\Sigma,\mathrm{p}]}, c_\alpha : !^\omega[\alpha_{[\mathrm{p}]}])\sigma' &=& \Theta_{[\Sigma,\mathrm{p}]}\sigma, c_\alpha : !^\omega[s_{[\mathrm{p}]}]\sigma \\
\Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma &= \mathbf{x}_{[\mathrm{p}]} : \mathbf{T}(\Sigma,G,[\mathrm{p}],n)\sigma &=& \mathbf{x}_{[\mathrm{p}]} : \$_1^0\, s_{[\mathrm{p}]}\sigma \\
\Gamma_{[\Sigma',G',\mathrm{p},n]}\sigma' &= \mathbf{x}_{[\mathrm{p}]} : \mathbf{T}(\Sigma',G',[\mathrm{p}],n)\sigma' &=& \mathbf{x}_{[\mathrm{p}]} : s_{[\mathrm{p}]}\sigma
\end{aligned}
$$

We derive:

$$
\frac{
\dfrac{\Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash \mathbf{x}_{[\mathrm{p}]} : \$_1^0\, s_{[\mathrm{p}]}\sigma}{c_\alpha : !^\omega[s_{[\mathrm{p}]}]\sigma, \Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash_1 c_\alpha!\langle \mathbf{x}_{[\mathrm{p}]}\rangle}\;[\text{T-OUT*}]
\qquad
\dfrac{\dfrac{\vdots}{\Theta_{[\Sigma',\mathrm{p}]}\sigma', \Gamma_{[\Sigma',G',\mathrm{p},n]}\sigma' \vdash_1 \mathbf{P}(\Sigma',G',\mathrm{p})}\;\text{IND.HYP.}}{\Theta_{[\Sigma,\mathrm{p}]}\sigma, c_\alpha : \#^\omega[s_{[\mathrm{p}]}]\sigma \vdash_1 *c_\alpha?(\mathbf{x}_{[\mathrm{p}]}).\mathbf{P}(\Sigma',G',\mathrm{p})}\;[\text{T-IN*}]
}{
\dfrac{\Theta_{[\Sigma,\mathrm{p}]}\sigma, c_\alpha : \#^\omega[s_{[\mathrm{p}]}]\sigma, \Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash_1 c_\alpha!\langle \mathbf{x}_{[\mathrm{p}]}\rangle \mid *c_\alpha?(\mathbf{x}_{[\mathrm{p}]}).\mathbf{P}(\Sigma',G',\mathrm{p})}{\Theta_{[\Sigma,\mathrm{p}]}\sigma, \Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma \vdash_1 \mathbf{P}(\Sigma,G,\mathrm{p})}\;[\text{T-NEW}]
}\;[\text{T-PAR}]
$$

$\boxed{G = \mathtt{A} \to \mathtt{B}@\ell.G' \text{ and } \mathrm{p} = \mathtt{A}}$ We consider the case $\mathtt{A} \wr \mathtt{B} \in \Sigma(G')$, the other being analogous. Then, from Lemma D.5, we know that

$$
\begin{aligned}
\mathbf{T}(\Sigma,G',\mathtt{A},\mathtt{B},\ell)\sigma &= p[s]_k^h \\
\mathbf{T}(\Sigma,G',\mathtt{B},\mathtt{A},\ell)\sigma &= \overline{p}[s]_k^h
\end{aligned}
$$

for some $h, k \in \mathbb{N}$, $p \in \{?,!\}$ such that $h > 0$, and s. We have:

$$
\begin{aligned}
\mathbf{P}(\Sigma,G,\mathrm{p}) &= (\nu a)(x_{\mathtt{AB}}!\langle a\rangle \mid \mathtt{let}\ x_{\mathtt{AB}} = a\ \mathtt{in}\ \mathbf{P}(\Sigma,G',\mathrm{p})) \\
\Gamma_{[\Sigma,G,\mathrm{p},n]}\sigma &= \{x_{\mathrm{pq}} : \$^{\ell-n}\,\mathbf{T}(\Sigma,G',\mathrm{p},\mathrm{q},\ell)\sigma\}_{\mathrm{q}\neq\mathtt{B}}, x_{\mathtt{AB}} : !\,[\overline{p}[s]_k^h]_0^{\ell-n} \\
\Gamma_{[\Sigma,G',\mathrm{p},\ell]}\sigma &= \{x_{\mathrm{pq}} : \mathbf{T}(\Sigma,G',\mathrm{p},\mathrm{q},\ell)\sigma\}_{\mathrm{q}\neq\mathtt{B}}, x_{\mathtt{AB}} : p[s]_k^h
\end{aligned}
$$

We derive $\mathscr{A}$:

$$
\frac{a : \overline{p}[s]_{k+1}^{\ell-n+h} \vdash a : \overline{p}[s]_{k+1}^{\ell-n+h}}{x_{\mathtt{AB}} : !\,[\overline{p}[s]_k^h]_0^{\ell-n}, a : \overline{p}[s]_{k+1}^{\ell-n+h} \vdash_1 x_{\mathtt{AB}}!\langle a\rangle}\;[\text{T-OUT}]
$$

as well as $\mathscr{B}$:

$$\cfrac{\cfrac{\cfrac{\vdots}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G',\mathbf{p},\ell]}\sigma \vdash_1 \mathbf{P}(\Sigma, G', \mathbf{p})} \text{ IND.HYP.}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \{x_{\mathbf{pq}} : \$^{\ell-n} \mathbf{T}(\Sigma, G, \mathbf{p}, \mathbf{q}, \ell)\sigma\}_{\mathbf{q}\neq\mathbf{B}}, x_{\mathbf{AB}} : p[s]_k^{\ell-n+h} \vdash_1 \mathbf{P}(\Sigma, G', \mathbf{p})} \text{ [T-LIFT]}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \{x_{\mathbf{pq}} : \$^{\ell-n} \mathbf{T}(\Sigma, G, \mathbf{p}, \mathbf{q}, \sigma)\}_{\mathbf{q}\neq\mathbf{B}}, a : p[s]_k^{\ell-n+h} \vdash_1 \texttt{let } x_{\mathbf{AB}} = a \texttt{ in } \mathbf{P}(\Sigma, G', \mathbf{p})} \text{ [T-LET]}$$

From these we obtain:

$$\cfrac{\cfrac{\mathscr{A} \qquad\qquad \mathscr{B}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G,\mathbf{p},n]}\sigma, a : \#[s]_{2k+1}^{\ell-n+h} \vdash_1 x_{\mathbf{AB}}!\langle a \rangle \mid \texttt{let } x_{\mathbf{AB}} = a \texttt{ in } \mathbf{P}(\Sigma, G', \mathbf{p})} \text{ [T-PAR]}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G,\mathbf{p},n]}\sigma \vdash_1 (\nu a)(x_{\mathbf{AB}}!\langle a \rangle \mid \texttt{let } x_{\mathbf{AB}} = a \texttt{ in } \mathbf{P}(\Sigma, G', \mathbf{p}))} \text{ [T-NEW]}$$

$\boxed{G = \mathbf{A} \to \mathbf{B}@\ell.G' \text{ and } \mathbf{p} = \mathbf{B}}$ We have:

$$
\begin{array}{rcl}
\mathbf{P}(\Sigma, G, \mathbf{p}) &=& x_{\mathbf{BA}}?(x_{\mathbf{BA}}).\mathbf{P}(\Sigma, G', \mathbf{p}) \\
\Gamma_{[\Sigma,G,\mathbf{p},n]}\sigma &=& \{x_{\mathbf{pq}} : \$^{\ell-n} \mathbf{T}(\Sigma, G', \mathbf{p}, \mathbf{q}, \ell)\sigma\}_{\mathbf{q}\neq\mathbf{A}}, x_{\mathbf{BA}} : ?[\mathbf{T}(\Sigma, G', \mathbf{B}, \mathbf{A}, \ell)\sigma]_0^{\ell-n} \\
\Gamma_{[\Sigma,G',\mathbf{p},\ell]}\sigma &=& \mathbf{x}_{[\mathbf{p}]} : \mathbf{T}(\Sigma, G', [\mathbf{p}], \ell)\sigma
\end{array}
$$

We derive:

$$\cfrac{\cfrac{\cfrac{\vdots}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G',\mathbf{p},\ell]}\sigma \vdash_1 \mathbf{P}(\Sigma, G', \mathbf{p})} \text{ IND.HYP.}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \{x_{\mathbf{pq}} : \mathbf{T}(\Sigma, G', \mathbf{p}, \mathbf{q}, \ell)\sigma\}_{\mathbf{q}\neq\mathbf{A}}, x_{\mathbf{BA}} : ?[\mathbf{T}(\Sigma, G', \mathbf{B}, \mathbf{A}, \ell)\sigma]_0^0 \vdash_1 x_{\mathbf{BA}}?(x_{\mathbf{BA}}).\mathbf{P}(\Sigma, G', \mathbf{p})} \text{ [T-IN]}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G,\mathbf{p},n]}\sigma \vdash_1 \mathbf{P}(\Sigma, G, \mathbf{p})} \text{ [T-LIFT]}$$

where the application of [T-IN] uses Lemma D.5 for deducing that the priorities of all the linear types in the environment are strictly positive.

$\boxed{G = \mathbf{A} \to \mathbf{B}@\ell.G' \text{ and } \mathbf{p} \notin \mathbf{A} \wr \mathbf{B}}$ We have:

$$
\begin{array}{rcl}
\mathbf{P}(\Sigma, G, \mathbf{p}) &=& \mathbf{P}(\Sigma, G', \mathbf{p}) \\
\Gamma_{[\Sigma,G,\mathbf{p},n]}\sigma &=& \mathbf{x}_{[\mathbf{p}]} : \mathbf{T}(\Sigma, G, [\mathbf{p}], n)\sigma = \mathbf{x}_{[\mathbf{p}]} : \$^{\ell-n} \mathbf{T}(\Sigma, G', [\mathbf{p}], \ell)\sigma = \$^{\ell-n} \Gamma_{[\Sigma,G',\mathbf{p},\ell]}\sigma
\end{array}
$$

We derive:

$$\cfrac{\cfrac{\vdots}{\Theta_{[\Sigma',\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G',\mathbf{p},\ell]}\sigma \vdash_1 \mathbf{P}(\Sigma, G', \mathbf{p})} \text{ IND.HYP.}}{\Theta_{[\Sigma,\mathbf{p}]}\sigma, \Gamma_{[\Sigma,G,\mathbf{p},n]}\sigma \vdash_1 \mathbf{P}(\Sigma, G, \mathbf{p})} \text{ [T-LIFT]}$$

$\square$

Theorem D.7 (Theorem C.3). *Let G be realizable. Then $\{x_{\mathbf{pq}} : \mathbf{T}(\emptyset, G, \mathbf{p}, \mathbf{q}, 0) \mid \mathbf{q} \in \mathscr{R} \setminus \{\mathbf{p}\}\} \vdash_1 \mathbf{P}(\emptyset, G, \mathbf{p})$ for every $\mathbf{p} \in \mathscr{R}$.*

Proof. Consequence of Lemma D.6 by taking $\sigma = \emptyset$. $\square$