Buffered Communication Analysis in Distributed Multiparty Sessions*

Pierre-Malo Deniélou and Nobuko Yoshida

Department of Computing, Imperial College London

Abstract. Many communication-centred systems today rely on asynchronous messaging among distributed peers to make efficient use of parallel execution and resource access. With such asynchrony, the communication buffers can happen to grow inconsiderately over time. This paper proposes a static verification methodology based on multiparty session types which can efficiently compute the upper bounds on buffer sizes. Our analysis relies on a uniform causality audit of the entire collaboration pattern — an examination that is not always possible from each end-point type. We extend this method to design algorithms that allocate communication channels in order to optimise the memory requirements of session executions. From these analyses, we propose two refinements methods which respect buffer bounds: a *global protocol refinement* that automatically inserts confirmation messages to guarantee stipulated buffer sizes and a *local protocol refinement* to optimise asynchronous messaging without buffer overflow. Finally our work is applied to overcome a buffer overflow problem of the multi-buffering algorithm.

1 Introduction

Session types for buffer bound analysis. The expensive cost of synchronous communications has led programmers to rely on asynchronous messaging for efficient network interactions. The downside is that non-blocking IO requires buffers that can grow inconsiderately over time, bringing systems to stop. The analysis and debugging of this phenomenon is mainly done by a tedious monitoring of the communicated messages of the whole distributed system. This paper shows that, when a global interaction pattern is explicitly specified as a *multiparty session* [1, 10, 14, 21], types can provide an effective way to statically verify buffer usage and communication optimisations, automatically guaranteeing safe and deadlock-free runs.

Session types, first introduced in [9, 19], can specify communication protocols by describing the sequences and types of read, write and choices on a given channel. For example, type $T_0 = !\langle \mathsf{nat} \rangle; !\langle \mathsf{string} \rangle; ?\langle \mathsf{real} \rangle; \mathsf{end}$, in the original binary session type syntax, expresses that a nat-value and string-value will be sent in that order, then that a real-value is expected as an input, and finally that the protocol ends.

We can use session types to calculate the upper bounds of the buffer sizes of asynchronous channels (message passing is non-blocking and order-preserving using FIFO buffers). For example, from type T_0, we can compute that the maximum number of messages that might be stored in a communication buffer is two, while a different type

* The work is partially supported by EPSRC EP/F003757/01 and EPSRC G015635/01.

$T_1 = !\langle \mathsf{nat} \rangle; ?\langle \mathsf{real} \rangle; !\langle \mathsf{string} \rangle; \mathsf{end}$ guarantees a maximum size of one, since the dual process engaged with T_1 is forced to consume a nat-value before sending the next real-message. This use of session types is informally observed in [6] and formally studied in [7] for binary session types. However, the binary case does not yield a direct extension to multiparty interactions as explained below.

Buffer bounds analysis in multiparty sessions. We start by illustrating the difficulties of such an analysis on a simple three party interaction (Example (a) below), where $s!\langle V \rangle$ is an output of V to s, $s?(x);P$ an input at s, and $\mu X.P$ a recursive agent:

Example (a)

(A) Alice=$\mu X.s_1!\langle 1 \rangle; s_3?(x); X$

(B) Bob =$\mu X.s_1?(x); s_2!\langle Orange \rangle; X$

(C) Carol=$\mu X.s_2?(x); s_3!\langle 2.4 \rangle; X$

Example (b)

(B_1) $\text{Bob}_1 = \mu X.s_1?(x); X$

(C_1) $\text{Carol}_1 = \mu X.s_3!\langle 2.4 \rangle; X$

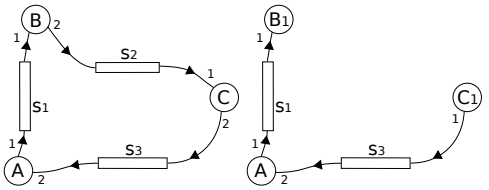

Example (a) **Example (b)**

We assume the three buffers of s_1, s_2 and s_3 are initially empty and that values are pushed and read one by one. Assuming session types ensure that accesses to buffers do not create any race condition at any moment of the infinite protocol execution, none of the channels s_1, s_2, s_3 need to buffer more that one value at any given time.

However, if we change Bob and Carol to Bob_1 and Carol_1 as Example (b) above, while they still interact correctly, the buffers of s_1 and s_3 need an unbounded size because of the *lack of synchronisation* between Bob_1 and Carol_1.

The main difficulty of the communication buffer analysis is that, unlike in binary session types, each end-point type itself does not provide enough information: for example, Alice's local type $T_a = \mu \mathbf{x}.s_1!\langle \mathsf{nat} \rangle; s_3?\langle \mathsf{real} \rangle; \mathbf{x}$ (repeatedly sends a nat-value to s_1 and receives a real-value from s_3) is *the same* in both Examples (a) and (b), while the needed buffer size for s_1 and s_3 are different (1 in (a) and ∞ in (b)) due to the change in the other parties' behaviours. Our first question is: *can we statically and efficiently determine the upper size of buffers in multiparty interactions?* In our case, we take advantage of the existence of a global session type [1, 10, 14, 21] for the analysis:

$$G = \mu \mathbf{x}.\text{Alice} \to \text{Bob}: s_1\langle \mathsf{nat} \rangle; \text{Bob} \to \text{Carol}: s_2\langle \mathsf{string} \rangle; \text{Carol} \to \text{Alice}: s_3\langle \mathsf{real} \rangle; \mathbf{x}$$

The above type represents the global interaction between Alice-Bob-Carol in (a) where $\text{Alice} \to \text{Bob}: s_1\langle \mathsf{nat} \rangle;$ means that Alice sends a nat-value to Bob through buffer s_1. To analyse buffer usage, we consider sessions as graphs and track *causal chains* for each channel: alternated message production and consumption mark the execution points at which buffers are emptied. This can be observed in Example (a). On the other hand, the global type of Alice-Bob_1-Carol_1 in (b) lacks the second $\text{Bob} \to \text{Carol}$: no message forces Carol to wait for Bob's reception before sending the next message. In that case, each buffer may accumulate an unbounded number of messages.

Channel allocation. Our next problem is about resource allocation. *Given a global scenario, can we assign the minimum number of resources (channels) without conflict* so that, for instance, we can efficiently open a minimal number of sockets for a given network interaction? Assume Alice and Carol in (a) wish to communicate one more message after completing three communications, where the new communication happens on a fresh channel s_4 (Example (c) below). Can we reuse either s_1, s_2 or s_3 for this new communication? Reusing s_1 creates a writing conflict (the order between Alice's

first and third messages would be unspecified) and reusing s_3 would create a reading conflict (Carol could read her own message).

Example (c)

(A_2) $\text{Alice}_2 = \mu X.s_1!\langle 1\rangle; s_3?(x); s_4!\langle x{+}1\rangle; X$

(B) $\text{Bob} = \mu X.s_1?(x); s_2!\langle \mathit{Orange}\rangle; X$

(C_2) $\text{Carol}_2 = \mu X.s_2?(x); s_3!\langle 2.4\rangle; s_4?(y); X$

Example (d)

(A_3) $\text{Alice}_3 = \mu X.s_1!\langle 1\rangle; s_3?(x); s_2!\langle x{+}1\rangle; X$

(C_3) $\text{Carol}_3 = \mu X.s_2?(x); s_3!\langle 2.4\rangle; s_2?(y); X$

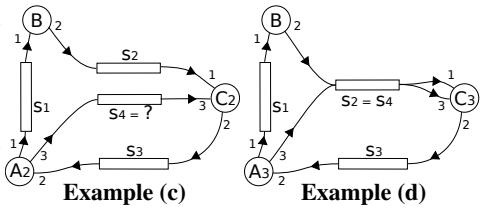

The only safe way to reuse a channel in Example (c) is to merge s_2 and s_4 as in Example (d), in which case communications on other channels prevent any conflict.

Global refinement for multiparty sessions. The third issue is how to fix a buffer overflow problem by "global refinement", i.e. alteration of the original global protocol to satisfy given buffer sizes. Here, our simple approach is the insertion of a minimal number of *confirmation messages* to enforce synchronisation. In network or business protocols, they can be implemented as a system level signal. Consider the interaction (b) among Alice-Bob_1-Carol_1 where each buffer requires an unbounded size. If we wish to enforce a buffer size of at most 2, we can build a new global type where one confirmation message from Bob to Carol is inserted in any second iteration as:

$$G' = \mu \mathbf{x}.\ \text{Alice} \to \text{Bob}\colon s_1\,\langle \mathsf{nat}\rangle; \text{Carol} \to \text{Alice}\colon s_3\,\langle \mathsf{real}\rangle; \\ \text{Alice} \to \text{Bob}\colon s_1\,\langle \mathsf{nat}\rangle; \text{Bob} \to \text{Carol}\colon s_2\,\langle \mathsf{string}\rangle; \text{Carol} \to \text{Alice}\colon s_3\,\langle \mathsf{real}\rangle; \mathbf{x}$$

The revised processes following G' are given as:

$$\text{Bob}_4 = \mu X.s_1?(x); s_1?(x); s_2!\langle \mathit{Signal}\rangle; X \qquad \text{Carol}_4 = \mu X.s_3!\langle 2.4\rangle; s_2?(x).s_3!\langle 2.4\rangle; X$$

Local refinement for multiparty messaging optimisations. The last issue is about flexible local refinement (optimisations) based on [13, 14]. Assume that, in Example (a), Bob wishes to always start the asynchronous transmission of the string *Orange* to the buffer s_3 without waiting for the delivery of the first nat-value from Alice on s_1.

$$\text{Bob}_5 = \mu X.s_2!\langle \mathit{Orange}\rangle; s_1?(x); X \tag{1.1}$$

Due to Bob's unilateral implementation change, all three minimal buffer sizes go up from 1 to 2. Moreover, suppose Bob repeatedly applies the same optimisation on his next n messages, as in $s_2!\langle \mathit{Orange}\rangle; s_2!\langle \mathit{Orange}\rangle; ..; s_2!\langle \mathit{Orange}\rangle;$Bob. While the result is communication-safe (no mismatch of the communication with Carol), all three minimal buffer sizes go up from 1 to n. *How can we perform local optimisation without altering buffer sizes in a multiparty session*?

Contributions are summarised in the figure below. To the best of our knowledge, our work is the first which guarantees safe buffered multiparty communications for the π-calculus with communication-safety, progress and flexible refinements. The key contribution is a general causal analysis over graphs constructed from multiparty session types (§ 3). The appendices list omitted definitions, examples and proofs.

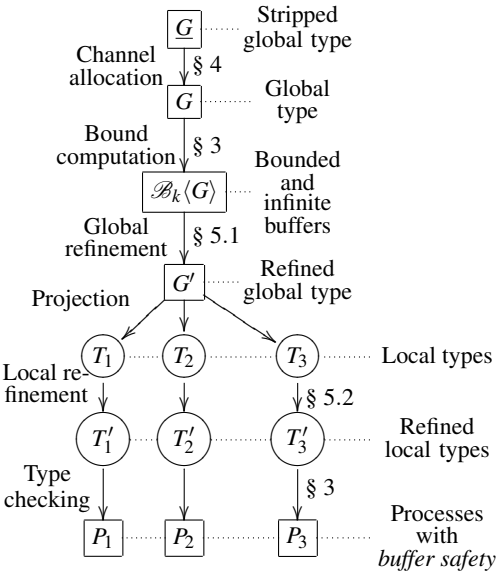

1. The overall analysis starts from global types that have no channel annotation. We attribute channels based on memory requirements (§ 4).
2. From global types, our bound analysis computes the buffer bounds of finite channels and finds the infinite ones (§ 3)
3. The global refinement method then introduces additional messages to prevent any unboundedness (§ 5.1).
4. Once the global type has been projected to local types, local refinement can optimise the distributed execution of the participants' processes (§ 5.2).
5. The running optimised processes can then be typed and enjoy communication, buffer and type safety and progress properties (§ 3).
6. We apply our work to the multibuffering algorithm and to a Multiprocessor System-on-Chip (MPSoC) use case (§ 6).

2 Asynchronous Multiparty Sessions

Syntax. We start from the π-calculus for multiparty sessions from [10] with unbounded and bounded buffers. Base sets and the grammars are given below.

P	$::= \overline{a}_{[2..\mathtt{n}]}(\tilde{s}^{\tilde{m}}).P \mid a_{[\mathtt{p}]}(\tilde{s}).P$	request, accept	$a,b,x,y,..$ shared names
	$\mid s!\langle\tilde{e}\rangle;P \mid s?(\tilde{x});P$	send, receive	$s,t,..$ session channels
	$\mid s!\langle\!\langle\tilde{s}\rangle\!\rangle;P \mid s?(\!(\tilde{s})\!);P$	session send, receive	$l,l',..$ labels
	$\mid s \lhd l;P \mid s \rhd \{l_i\colon P_i\}_{i\in I}$	selection, branch	$X,Y,..$ process variables
	$\mid$ if e then P else Q	conditional	$m,n,..$ buffer size (integers or ∞)
	$\mid \mathbf{0} \mid (\nu\, a)P \mid (\nu\, \tilde{s})P$	inact, hiding	$e ::= v \mid e$ and $e' \cdots$ expressions
	$\mid P \mid Q \mid \mu X.P \mid X$	par, recursion	$v ::= a \mid$ true $\mid$ false $\cdots$ values
	$\mid s^n : \tilde{h}$	message buffer	$h ::= l \mid \tilde{v} \mid \tilde{t}$ message values

$\overline{a}_{[2..\mathtt{n}]}(\tilde{s}^{\tilde{m}}).P$ initiates, through a shared name a, a new session s_i with buffer size m_i ($1 \le n \le \infty$) with other participants, each of the form $a_{[\mathtt{p}]}(\tilde{s}).Q$ with $1 \le \mathtt{p} \le n-1$. The s_i in vector $\tilde{s}$ is a session channel (bounded by buffer size m_i) used in the session. We call p, q,... (natural numbers) the *participants* of a session. Session communications (which take place inside an established session) are performed by the sending and receiving of a value; the session sending and reception (where the former delegates to the latter the capability to participate in a session by passing a channel associated with the session which is called *delegation*); and by selection and branching (the former chooses one of the branches offered by the latter). $s^n : \tilde{h}$ is a *message buffer of size n* representing ordered messages in transit $\tilde{h}$ with destination s. This may be considered as a network pipe in a TCP-like transport with fixed bandwidth. The rest of the syntax is standard from [10]. We often omit n from $s^n : \tilde{h}$, $\mathbf{0}$, and unimportant arguments e.g. $s!\langle\rangle$ and $s?();P$. An *initial* process does not contain any runtime syntax (buffers and session hiding).

Reductions. A selection of reduction rules is given below.

$$\overline{a}_{[2..\mathtt{n}]}(\tilde{s}^{\tilde{n}}).P_1 \mid a_{[2]}(\tilde{s}).P_2 \mid \cdots \mid a_{[\mathtt{n}]}(\tilde{s}).P_{\mathtt{n}} \;\to\; (\nu\, \tilde{s})(P_1 \mid P_2 \mid \ldots \mid P_{\mathtt{n}} \mid s_1^{n_1}:\emptyset \mid \ldots \mid s_m^{n_m}:\emptyset)$$

$$s!\langle \tilde{e}\rangle;P \mid s^n:\tilde{h} \;\to\; P \mid s^n:\tilde{h}\cdot\tilde{v} \;\; (n \gtrsim |\tilde{h}|,\; e_i \downarrow v_i) \qquad s?(\tilde{x});P \mid s^n:\tilde{v}\cdot\tilde{h} \to P[\tilde{v}/\tilde{x}] \mid s^n:\tilde{h}$$

$$s!\langle\!\langle \tilde{t}\rangle\!\rangle;P \mid s^n:\tilde{h} \;\to\; P \mid s^n:\tilde{h}\cdot\tilde{t} \;\; (n \gtrsim |\tilde{h}|) \qquad s?(\!(\tilde{t})\!);P \mid s^n:\tilde{t}\cdot\tilde{h} \to P \mid s^n:\tilde{h}$$

$$s \lhd l;P \mid s^n:\tilde{h} \;\to\; P \mid s^n:\tilde{h}\cdot l \;\; (n \gtrsim |\tilde{h}|) \qquad s \rhd \{l_i\colon P_i\}_{i\in I} \mid s^n:l_j\cdot\tilde{h} \to P_j \mid s^n:\tilde{h} \quad (j\in I)$$

The first rule describes the initiation of a new session among $\mathtt{n}$ participants that synchronise over the shared name a. After the initiation, they will share m fresh private session channels s_i and the associated m empty buffers of size n_m ($\emptyset$ denotes an empty queue). The output rules for values, sessions and selection respectively enqueue values, sessions and labels if the buffer is not full. $e_i \downarrow v_i$ denotes the evaluation of e_i to v_i. We define $|\emptyset| = 0$ and $|\tilde{h}\cdot h| = |\tilde{h}| + 1$. The size $n = \infty$ corresponds to the original asynchronous unbounded buffered semantics [10]. The input rules perform the complementary operations. Processes are considered modulo a structural equivalence $\equiv$, whose definition is standard (e.g. $\mu X.P \equiv P[\mu X.P/X]$) [10].

3 Bound Analysis in Multiparty Sessions

This section presents an analysis of causal chains and buffer sizes and introduces the typing system for the *buffer safety* property (Corollary 3.9).

3.1 Global Types and Dependencies

Global types. A *global type*, written by $G, G', ..$, describes the whole conversation scenario of a multiparty session as a type signature. Our starting syntax is from [10].

G,G'	$::=$ $\mathtt{p}\to\mathtt{p}'\colon k\langle U\rangle;G'$	values	U,U'	$::=$ $\tilde{S} \mid T@\mathtt{p}$	sorts, session
	$\mid$ $\mathtt{p}\to\mathtt{p}'\colon k\{l_j\colon G_j\}_{j\in J}$	branching	S,S'	$::=$ $\mathsf{bool} \mid \mathsf{nat} \mid G$	base, shared
	$\mid$ $\mu\mathbf{x}.G \mid \mathbf{x} \mid \mathsf{end}$	recursion, end			

Type $\mathtt{p}\to\mathtt{p}'\colon k\langle U\rangle;G'$ says that participant $\mathtt{p}$ sends a message of type U on the channel k (represented as a natural number) so that participant $\mathtt{p}'$ can receive it. The session continues with the interactions described in G'. The *value types* U,U' are either a vector of sorts or a *located type* $T@\mathtt{p}$, representing a local type T assigned to participant $\mathtt{p}$. Located types are used for delegation and defined in § 3.3. *Sorts* S,S' are either base types or global types for shared names. Type $\mathtt{p}\to\mathtt{p}'\colon k\{l_j\colon G_j\}_{j\in J}$ says that participant $\mathtt{p}$ can invoke one of the l_i labels on channel k (for participant $\mathtt{p}'$ to read) and that interactions described in G_j follow. We require $\mathtt{p}\neq\mathtt{p}'$ to prevent self-sent messages. Type $\mu\mathbf{x}.G$ is for recursive protocols, assuming the type variables $(\mathbf{x},\mathbf{x}',\ldots)$ are guarded in the standard way, i.e. they only occur under values or branchings. We assume G in value types is closed, i.e. without free type variables. Type end represents session termination (often omitted). $k\in G$ means k appears in G. The functions $chans(G)$ and $prins(G)$ respectively give the number of channels and participants of G.

Sessions as graphs. Global types can be seen (isomorphically) as *session graphs*, that we define in the following way. First, we annotate in G each syntax occurrence of

subterms of the form $\mathtt{p} \to \mathtt{p}' \colon k\,\langle U \rangle; G'$ or $\mathtt{p} \to \mathtt{p}' \colon k\,\{l_j \colon G_j\}_{j \in J}$ with a node name $(\mathtt{n}_1, \mathtt{n}_2, \ldots)$. Then, we inductively define a function $\mathtt{node}_G$ that gives a node $\mathtt{n}_k$ (or the special node end) for each of the syntactic subterm of G as follows:

$$\begin{array}{rlrl} \mathtt{node}_G(\mathsf{end}) &= \mathsf{end} & \mathtt{node}_G(\mathtt{n}_i : \mathtt{p} \to \mathtt{p}' \colon k\,\langle U \rangle; G') &= \mathtt{n}_i \\ \mathtt{node}_G(\mu \mathbf{x}.G') &= \mathtt{node}_G(G') & \mathtt{node}_G(\mathtt{n}_j : \mathtt{p} \to \mathtt{p}' \colon k\,\{l_j \colon G_j\}_{j \in J}) &= \mathtt{n}_j \\ \mathtt{node}_G(\mathbf{x}) &= \mathtt{node}_G(\mu \mathbf{x}.G') & \textit{(if the binder of } \mathbf{x} \textit{ is } \mu \mathbf{x}.G' \in G) & \end{array}$$

We define G as a session graph in the following way: for each subterm of G of the form $\mathtt{n} : \mathtt{p} \to \mathtt{p}' \colon k\,\langle U \rangle; G'$, we have an edge from $\mathtt{n}$ to $\mathtt{node}_G(G')$, and for each subterm of G of the form $\mathtt{n}' : \mathtt{p} \to \mathtt{p}' \colon k\,\{l_j \colon G_j\}_{j \in J}$, we have edges from $\mathtt{n}'$ to each of the $\mathtt{node}_G(G_j)$ for $j \in J$. We also define the functions $\mathtt{pfx}(\mathtt{n}_i)$ and $\mathtt{ch}(\mathtt{n}_i)$ that respectively give the prefix $(\mathtt{p} \to \mathtt{p}' : k)$ and channel (k) that correspond to $\mathtt{n}_i$. For a global type G, $\mathtt{node}_G(G)$ distinguishes the *initial* node. *size*(G) denotes the number of edges of G.

Example 3.1 (Session graph) Our running example extends Example (a) from § 1 with branching. Below, we give the global type followed by its graph representation, with the edges as the dotted arrows (labels are for information). $\mathtt{n}_1$ is the initial node.

$$\mu \mathbf{x}.\mathrm{Alice} \to \mathrm{Bob} \colon s_1\,\langle \mathsf{nat} \rangle; \mathrm{Bob} \to \mathrm{Carol} : s_2 \left\{ \begin{array}{l} l_1 : \mathrm{Carol} \to \mathrm{Alice} : s_3 \{ l_3 : \mathrm{Alice} \to \mathrm{Carol} \colon s_2\,\langle \mathsf{string} \rangle; \mathbf{x} \} \\ l_2 : \mathrm{Carol} \to \mathrm{Alice} : s_3 \{ l_4 : \mathsf{end} \} \end{array} \right\}$$

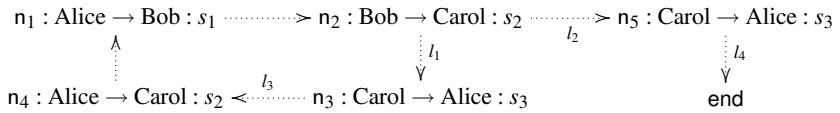

The recursion call yields a cycle in the graph, while branching gives the edges l_1 and l_2.

The edges of a given session graph G define a successor relation between nodes, written $\mathtt{n} \prec \mathtt{n}'$ (omitting G). Paths in this session graph are referred to by the sequence of nodes they pass through: a path $\mathtt{n}_0 \prec \ldots \prec \mathtt{n}_n$ can be written more concisely $\mathtt{n}_0 \ldots \mathtt{n}_n$ or $\tilde{\mathtt{n}}$ when there is no ambiguity. We say that a path $\mathtt{n}_0 \ldots \mathtt{n}_n$ *has suffix* $\mathtt{n}_i \ldots \mathtt{n}_n$ for $0 < i < n$. The empty path is ε. The transitive closure of $\prec$ is $\prec\!\!\!\prec$.

IO-chains. We detect causality chains in a given G by the relation $\prec_{\mathtt{IO}}$, defined below:

$$\mathtt{n}_1 \prec_{\mathtt{IO}} \mathtt{n}_2 \;\text{ if } \mathtt{n}_1 \prec\!\!\!\prec \mathtt{n}_2 \text{ and } \mathtt{pfx}(\mathtt{n}_1) = \mathtt{p}_1 \to \mathtt{p} : k_1 \text{ and } \mathtt{pfx}(\mathtt{n}_2) = \mathtt{p} \to \mathtt{p}_2 : k_2 \text{ with } k_1 \neq k_2$$

The relation $\prec_{\mathtt{IO}}$ asserts the order between a reception by a principal and the next message it sends. An *input-output dependency (IO-dependency)* from $\mathtt{n}_1$ to $\mathtt{n}_n$ is a chain $\mathtt{n}_1 \prec_{\mathtt{IO}} \cdots \prec_{\mathtt{IO}} \mathtt{n}_n$ $(n \geq 1)$.

3.2 Algorithms for Buffer Size Analysis

Unbounded buffers. In some sessions, messages (sent asynchronously) can accumulate without moderation in a buffer. A simple test can predict which channels require an unbounded buffer. We use the fact that IO-dependencies characterise the necessity for a channel buffer to be emptied before proceeding. Infinite channels are the ones where such a dependency is missing.

Definition 3.2 (infinite and finite) A channel k is said to be *finite* if, for every node $\mathsf{n} \in G$ and for every cycle $\tilde{\mathsf{n}}$ for $\prec$ such that $\mathtt{ch}(\mathsf{n}) = k$ and $\mathsf{n} \in \tilde{\mathsf{n}}$, there exists a cycle for $\prec_{\mathtt{IO}}$ that starts from n and only involves nodes from $\tilde{\mathsf{n}}$. The other channels are *infinite*.

Correspondingly, buffers are said to be *bounded* or *unbounded*. Given G, checking for the infinity of k in G can be computed in a time bounded by $O(size(G)^3)$. The proof relies on the fact that establishing all IO-dependencies of a given session has $O(size(G)^3)$ time-complexity (assuming the number of participants as a constant).

Example 3.3 (Session graph and infinite channels) We illustrate on our running example the previous notions. We add to the picture the IO-dependencies (with $\Rightarrow$).

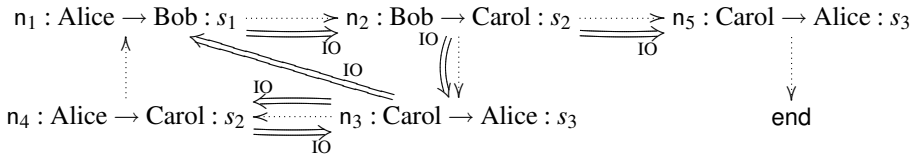

Since each node of the main cycle $\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3\mathsf{n}_4$ is part of the IO-cycles $\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3$ or $\mathsf{n}_3\mathsf{n}_4$, there are no infinite channels in this session.

Counting finite buffer size. To compute the bounds on buffer sizes, we first need to define a property on paths that characterises when a buffer has to be emptied.

Definition 3.4 (reset) If $\tilde{\mathsf{n}} = \mathsf{n}_0 \ldots \mathsf{n}_n\mathsf{n}$ is a path in G, the property $\mathrm{Reset}(\tilde{\mathsf{n}})$ holds if there exist $0 \leq i_0 < \ldots < i_j \leq n$ $(j \geq 1)$ such that $\mathsf{n}_{i_0} \prec_{\mathtt{IO}} \ldots \prec_{\mathtt{IO}} \mathsf{n}_{i_j} \prec_{\mathtt{IO}} \mathsf{n}$ and $\mathtt{ch}(\mathsf{n}_{i_0}) = \mathtt{ch}(\mathsf{n})$. One practical instance of the nodes $\{\mathsf{n}_{i_0}, \ldots, \mathsf{n}_{i_j}, \mathsf{n}\}$ is called the reset nodes of $\tilde{\mathsf{n}}$.

The paths that satisfy the reset property are the ones for which there exists a reception guard to the last node.

Now that we know which buffers are infinite and have characterised the resetting paths that control buffer growth, we can describe our algorithm to count the buffer size required by finite channels. For each channel k of a global session type G, we define a function $\mathscr{B}_k\langle G\rangle$ that will compute the bound on the buffer size of channel k. The key step is to reset the counter when we recognise the appropriate IO-dependencies.

Definition 3.5 (bound computation) Given a session graph G, for each channel k, we compute the bound as $\mathscr{B}_k\langle G\rangle = \mathscr{B}_k\langle 0, \emptyset, \varepsilon, \mathsf{n}_0\rangle$ for n_0 the initial node of G.

$$\mathscr{B}_k\langle m, \mathscr{P}, \tilde{\mathsf{n}}, \mathsf{n}\rangle = \begin{cases} 0 & \text{if } \mathsf{n} = \mathsf{end} \text{ or } \tilde{\mathsf{n}} \in \mathscr{P} \\ \max_{\mathsf{n} \prec \mathsf{n}'} \mathscr{B}_k\langle m, \{\tilde{\mathsf{n}}\} \cup \mathscr{P}, \tilde{\mathsf{n}}\mathsf{n}, \mathsf{n}'\rangle & \text{if } \mathtt{ch}(\mathsf{n}) = k', k \neq k' \\ \max_{\mathsf{n} \prec \mathsf{n}'} \mathscr{B}_k\langle 1, \{\tilde{\mathsf{n}}\} \cup \mathscr{P}, \mathsf{n}, \mathsf{n}'\rangle & \text{if } \mathtt{ch}(\mathsf{n}) = k, \mathrm{Reset}(\tilde{\mathsf{n}}\mathsf{n}) \\ \max(m+1, \max_{\mathsf{n} \prec \mathsf{n}'} \mathscr{B}_k\langle m+1, \{\tilde{\mathsf{n}}\} \cup \mathscr{P}, \tilde{\mathsf{n}}\mathsf{n}, \mathsf{n}'\rangle) & \text{if } \mathtt{ch}(\mathsf{n}) = k, \neg\mathrm{Reset}(\tilde{\mathsf{n}}\mathsf{n}) \end{cases}$$

The algorithm explores all the paths of the session graph until they grow to satisfy the reset property. Since we examine only finite channels, the length of such paths is limited and the algorithm terminates. The bound on the buffer size of a channel is the maximum buffer size required over these paths. For each path, the algorithm acts recursively on the edges and maintains a counter (m in $\mathscr{B}_k\langle m, \mathscr{P}, \tilde{\mathsf{n}}, \mathsf{n}\rangle$) that records the current size of the buffer. If the current prefix does not involve the channel k, the size of the buffer

is unchanged and the computation continues to the next nodes. If the current prefix uses the channel k, there are two cases: (a) the reset property holds for the current path, in which case the buffer has been emptied prior to the current message; or (b) the reset property does not hold and the buffer needs to be able to keep one more value. When there are no further node, or when the path currently examined has already been considered (i.e. is in $\mathscr{P}$), the algorithm stops.

Given a global type G, the upper bound of channel k in G can be computed in polynomial time. Note that the computation can be done for all channels at once.

Example 3.6 (buffer bound analysis) We illustrate the algorithm on our running session example, where we compute the bound for channel s_2 (we omit $\mathscr{P}$ for readability):

	max	explanation
$\mathscr{B}_{s_2}\langle 0, \varepsilon, \mathsf{n}_1\rangle$		
$= \mathscr{B}_{s_2}\langle 0, \mathsf{n}_1, \mathsf{n}_2\rangle$	0	$s_1 \neq s_2$
$= \max(\mathscr{B}_{s_2}\langle 1, \mathsf{n}_1\mathsf{n}_2, \mathsf{n}_3\rangle, \mathscr{B}_{s_2}\langle 1, \mathsf{n}_1\mathsf{n}_2, \mathsf{n}_5\rangle)$	1	$\neg\text{Reset}(\mathsf{n}_1\mathsf{n}_2)$
$= \max(\mathscr{B}_{s_2}\langle 1, \mathsf{n}_1\mathsf{n}_2\mathsf{n}_3, \mathsf{n}_4\rangle, \mathscr{B}_{s_2}\langle 1, \mathsf{n}_1\mathsf{n}_2\mathsf{n}_5, \mathsf{end}\rangle)$	1	$s_3 \neq s_2$
$= \max(\mathscr{B}_{s_2}\langle 1, \mathsf{n}_4, \mathsf{n}_1\rangle, 0)$	1	$\text{Reset}(\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3\mathsf{n}_4)$
$= \mathscr{B}_{s_2}\langle 1, \mathsf{n}_4\mathsf{n}_1, \mathsf{n}_2\rangle$	1	$s_1 \neq s_2$
$= \max(\mathscr{B}_{s_2}\langle 2, \mathsf{n}_4\mathsf{n}_1\mathsf{n}_2, \mathsf{n}_3\rangle, \mathscr{B}_{s_2}\langle 2, \mathsf{n}_4\mathsf{n}_1\mathsf{n}_2, \mathsf{n}_5\rangle)$	2	$\neg\text{Reset}(\mathsf{n}_4\mathsf{n}_1\mathsf{n}_2)$
$= \max(\mathscr{B}_{s_2}\langle 2, \mathsf{n}_4\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3, \mathsf{n}_4\rangle, \mathscr{B}_{s_2}\langle 2, \mathsf{n}_4\mathsf{n}_1\mathsf{n}_2\mathsf{n}_5, \mathsf{end}\rangle)$	2	$s_3 \neq s_2$
$= \max(\mathscr{B}_{s_2}\langle 1, \mathsf{n}_4, \mathsf{n}_1\rangle, 0)$	**2**	$\text{Reset}(\mathsf{n}_4\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3\mathsf{n}_4)$

The algorithm starts with n_1, the root of G. Since n_1 uses buffer s_1 (different from s_2), we continue with the successor n_2. It uses s_2 and, since the accumulated path $\mathsf{n}_1\mathsf{n}_2$ does not satisfy the reset property, the buffer requirement of s_2 needs to be increased to 1. The next nodes, n_3 and n_5, do not use the channel s_2. Since n_4 uses s_2 and $\text{Reset}(\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3\mathsf{n}_4)$ holds (there is $\mathsf{n}_2 \prec_{\mathtt{IO}} \mathsf{n}_3 \prec_{\mathtt{IO}} \mathsf{n}_4$), the buffer has to be emptied before n_4: we thus reinitialise the buffer requirement to 1 and the path to just n_4. On the other branch, we reach end and stop the computation. The next prefix of n_4, n_1, does not use s_2, but it successor n_2 does. We thus check the reset property on the path $\mathsf{n}_4\mathsf{n}_1\mathsf{n}_2$, but it does not hold. The buffer requirement is thus increased to 2. As previously, n_3 and n_5 do not use the channel s_2 and the accumulated path (in the main branch) becomes $\mathsf{n}_4\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3$. The next prefix, n_4, uses s_2 and $\text{Reset}(\mathsf{n}_4\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3\mathsf{n}_4)$ holds: thus we initialise the buffer requirement back to 1 and the path to just n_4. However, we just explored such a situation earlier in the computation and thus stop. The maximum buffer size encountered for s_2 is then 2. Such a computation for s_1 and s_3 gives a buffer size of 1.

3.3 Subject Reduction and Buffer Safety

Once global type G is agreed upon by all parties, a local type T_i from each party's viewpoint is generated as a projection of G, and implemented as a process P_i. If all the resulting local processes $P_1, .., P_n$ can be type-checked against $T_1, .., T_n$, they are automatically guaranteed to interact properly, without communication mismatch (communication safety) nor getting stuck inside a session (progress) [10]. Here we additionally ensure the absence of buffer-overflow based on the buffer bound analysis of G.

Local types. Local session types type-abstract sessions from each end-point's view.

$$T ::= k!\langle U\rangle;T \mid k?\langle U\rangle;T \mid k \oplus \{l_i\colon T_i\}_{i\in I} \mid k\&\{l_i\colon T_i\}_{i\in I} \mid \mu \mathbf{x}.T \mid \mathbf{x} \mid \mathsf{end}$$

Type $k!\langle U\rangle$ expresses the sending to k of a value of type U. Type $k?\langle U\rangle$ is its dual. Type $k \oplus \{l_i\colon T_i\}_{i\in I}$ represents the transmission to k of a label l_i chosen in the set $\{l_i \mid i \in I\}$, followed by the communications described by T_i. Type $k\&\{l_i\colon T_i\}_{i\in I}$ is its dual. The remaining type constructors are standard. We say a type is *guarded* if it is neither a recursive type nor a type variable. The relation between global and local types is formalised by *projection*, written $G \upharpoonright \mathtt{p}$ (called *projection of G onto* $\mathtt{p}$) and defined in [10, 21]. For example, $(\mathtt{p} \to \mathtt{p}'\colon k\langle U\rangle; G') \upharpoonright \mathtt{p} = k!\langle U\rangle; (G' \upharpoonright \mathtt{p})$, $(\mathtt{p} \to \mathtt{p}'\colon k\langle U\rangle; G') \upharpoonright \mathtt{p}' = k?\langle U\rangle; (G' \upharpoonright \mathtt{p}')$ and $(\mathtt{p} \to \mathtt{p}'\colon k\langle U\rangle; G') \upharpoonright \mathtt{q} = (G' \upharpoonright \mathtt{q})$. We take an *equi-recursive* view, not distinguishing between $\mu\mathbf{x}.T$ and its unfolding $T[\mu\mathbf{x}.T/\mathbf{x}]$.

Linearity. To avoid race conditions and conflicts between typed processes, we build on the definition of linearity from [10]. The relations $\prec_{\mathtt{II}}$ and $\prec_{\mathtt{OO}}$ are defined by:

$\mathsf{n}_1 \prec_{\mathtt{II}} \mathsf{n}_2$ if $\mathsf{n}_1 \prec\!\!\prec \mathsf{n}_2$ and $\mathtt{pfx}(\mathsf{n}_1) = \mathtt{p}_1 \to \mathtt{p}\colon k_1$ and $\mathtt{pfx}(\mathsf{n}_2) = \mathtt{p}_2 \to \mathtt{p}\colon k_2$ s.t. $\mathtt{p}_1 \neq \mathtt{p}_2 \Leftrightarrow k_1 \neq k_2$

$\mathsf{n}_1 \prec_{\mathtt{OO}} \mathsf{n}_2$ if $\mathsf{n}_1 \prec\!\!\prec \mathsf{n}_2$ and $\mathtt{pfx}(\mathsf{n}_1) = \mathtt{p} \to \mathtt{p}_1\colon k_1$ and $\mathtt{pfx}(\mathsf{n}_2) = \mathtt{p} \to \mathtt{p}_2\colon k_2$ s.t. $\mathtt{p}_1 \neq \mathtt{p}_2 \Rightarrow k_1 \neq k_2$

The three relations $\prec_{\mathtt{IO}}$, $\prec_{\mathtt{II}}$ and $\prec_{\mathtt{OO}}$ are used to characterise the authorised sequences of actions. An *input dependency (I-dependency) from* n_1 *to* n_2 is a chain $\mathsf{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathsf{n}_2$ ($n \geq 1$) such that $\phi_i = \mathsf{IO}$ for $1 \leq i \leq n-1$ and $\phi_n = \mathsf{II}$. An *output dependency (O-dependency) from* n_1 *to* n_2 is a chain $\mathsf{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathsf{n}_2$ ($n \geq 1$) such that $\phi_i \in \{\mathsf{OO}, \mathsf{IO}\}$. These dependency relations are respectively written $\prec\!\!\prec_{\mathtt{II}}$ and $\prec\!\!\prec_{\mathtt{OO}}$. G is *linear* (written $\mathrm{Lin}(G)$) if, whenever two nodes $\mathsf{n}_1 \prec\!\!\prec \mathsf{n}_2$ use the same channel k, the dependencies $\mathsf{n}_1 \prec\!\!\prec_{\mathtt{II}} \mathsf{n}_2$ and $\mathsf{n}_1 \prec\!\!\prec_{\mathtt{OO}} \mathsf{n}_2$ hold. If G carries other global types, we inductively demand the same. Examples can be found in [10] and Appendix B.1 . We call linear global types whose projections are defined *coherent*. Hereafter we only consider coherent types.

Typing initial processes. The type judgements for initial processes are of the form $\Gamma \vdash P \triangleright \Delta$ which means: "under the environment Γ, process P has typing Δ". Environments are defined by: $\Gamma ::= \emptyset \mid \Gamma, u\colon S \mid \Gamma, X\colon \Delta$ and $\Delta ::= \emptyset \mid \Delta, \tilde{s}^{\tilde{m}}\colon \{T@\mathtt{p}\}_{\mathtt{p}\in I}$. A *sorting* $(\Gamma, \Gamma', ..)$ is a finite map from names to sorts and from process variables to sequences of sorts and types. *Typing* $(\Delta, \Delta', ..)$ records linear usage of session channels. In multiparty sessions, it assigns a family of located types to a vector of session channels. In addition, we annotate each session channel s_k with its buffer bound m_k.

Among the typing rules, the rule for session initiation uses the buffer size $\mathscr{B}_{s_i}\langle G\rangle$ calculated from G.

$$\frac{\Gamma \vdash a\colon G \quad \Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}}\colon (G \upharpoonright \mathtt{1})@\mathtt{1} \;\; |\tilde{s}| = \mathit{chans}(G) \;\; \mathscr{B}_k\langle G\rangle = m_k}{\Gamma \vdash \overline{a}_{[2..\mathtt{n}]}(\tilde{s}^{\tilde{m}}).P \triangleright \Delta}$$

The type for $\tilde{s}$ is the *first* projection of the declared global type for a in Γ. The end-point type $(G \upharpoonright \mathtt{p})@\mathtt{p}$ means that the participant $\mathtt{p}$ has $G \upharpoonright \mathtt{p}$, which is the projection of G onto $\mathtt{p}$, as its end-point type. The condition $|\tilde{s}| = \mathit{chans}(G)$ means the number of session channels meets those in G. The condition $\mathscr{B}_k\langle G\rangle = m_k$ ensures that the size of the buffer m_i for each s_k does not exceed the size calculated from G. Similarly for accept. Other rules for initial processes are identical with [10]. Note that since $\mathscr{B}_k\langle G\rangle$ is decidable, type-checking for processes with type annotations is decidable [10, 21].

The rest of the typing system for programs and one for runtime are similar with those in [10] (Appendix C.4). Judgements for runtime are there extended to $\Gamma \vdash_\Sigma P \triangleright \Delta$ with Σ a set of session channels associated to the current queue.

For the subject reduction, we need to keep track of the correspondence between the session environment and the buffer sizes. We use the reduction over session typing, $\Delta \xrightarrow{k} \Delta'$, that is generated by rules between types such as $k!\langle U\rangle;T@\mathtt{p}, k?\langle U\rangle;T'@\mathtt{q} \xrightarrow{k} T@\mathtt{p}, T'@\mathtt{q}$. The key lemma about the correspondence between buffer size and reduction follows. We set $[\![G]\!]$ to be the family $\{(G\upharpoonright \mathtt{p})@\mathtt{p} \mid \mathtt{p} \in G\}$. Regarding each type in $[\![G]\!]$ as the corresponding regular tree, we can define $\prec$, $\prec_{\mathtt{II}}$, $\prec_{\mathtt{IO}}$ and $\prec_{\mathtt{OO}}$ among its prefixes precisely as we have done for G.

Lemma 3.7 *If $\Delta(\tilde{s}) = [\![G]\!]$ and $\Delta \xrightarrow{s_k} \Delta'$, then $[\![G]\!](\xrightarrow{k})^*[\![G']\!]$ with $\Delta'(\tilde{s}) = [\![G']\!]$ and $\mathscr{B}_k\langle G\rangle \geq \mathscr{B}_k\langle G'\rangle$.*

When $\Gamma \vdash_\Sigma P \triangleright \Delta$, we say that $(\Gamma, \Sigma, P, \Delta)$ is *fully coherent* for session $\tilde{s}$ if there exist $P_1, \ldots, P_k, \Sigma', \Delta'$ such that $\Gamma \vdash_{\Sigma \uplus \Sigma'} P \mid P_1 \mid \ldots \mid P_k \triangleright \Delta, \Delta'$ and $\Delta, \Delta' = \Delta'', \tilde{s}^{\tilde{n}} : \{T_\mathtt{p}@\mathtt{p}\}_{\mathtt{p}\in I}$ with $[\![G]\!] = \{T_\mathtt{p}@\mathtt{p}\}_{\mathtt{p}\in I}$, G coherent and $\mathscr{B}_i\langle G\rangle \leq n_i$ $(1 \leq i \leq k)$.

Theorem 3.8 (Subject Reduction) *$\Gamma \vdash_\Sigma P \triangleright \Delta$ and $P \longrightarrow Q$ with $(\Gamma, \Sigma, P, \Delta)$ fully coherent imply $\Gamma \vdash_\Sigma Q \triangleright \Delta'$ for some Δ', s_k such that $\Delta = \Delta'$ or $\Delta(\xrightarrow{s_k})^*\Delta'$ and $\mathscr{B}_k\langle G\rangle \geq \mathscr{B}_k\langle G'\rangle$ with $\Delta(\tilde{s}) = [\![G]\!]$, $\Delta'(\tilde{s}) = [\![G']\!]$ and $(\Gamma, \Sigma, Q, \Delta')$ fully coherent.*

The proof relies on Lemma 3.7 and the fact that session reduction does not affect the causal dependencies within global types, so that buffer sizes can only decrease.

To state our buffer safety result, we define the *buffer overflow error* as follows:

$$\begin{array}{rcl} n \leq |\tilde{h}| & \Rightarrow & s!\langle \tilde{e}\rangle;P \mid s^n : \tilde{h} \rightarrow \mathsf{Err},\ s!\langle\!\langle \tilde{t}\rangle\!\rangle;P \mid s^n : \tilde{h} \rightarrow \mathsf{Err},\ s \lhd l;P \mid s^n : \tilde{h} \rightarrow \mathsf{Err} \\ P \rightarrow \mathsf{Err} & \Rightarrow & P \mid Q \rightarrow \mathsf{Err},\ (\nu\ a)P \rightarrow \mathsf{Err},\ (\nu\ \tilde{s})P \rightarrow \mathsf{Err},\ P \equiv Q \rightarrow \mathsf{Err} \end{array}$$

Corollary 3.9 (Buffer Safety) *If $\Gamma \vdash_\Sigma P \triangleright \Delta$, then for all P' s.t. $P \longrightarrow^* P'$, $P' \not\rightarrow \mathsf{Err}$.*

4 Channel Attribution

This section describes algorithms that attribute channels to the communications of a given global type without channels, called *stripped global types* $(\underline{G}, \underline{G}', \ldots)$ defined as:

$$\underline{G} ::= \ \ldots \mid \mathtt{p} \rightarrow \mathtt{p}'\langle U\rangle;\underline{G}' \mid \mathtt{p} \rightarrow \mathtt{p}'\{l_j : \underline{G}_j\}_{j\in J} \quad \text{values, branching}$$

Our algorithms transform $\underline{G}$ into regular type G by adding channel annotations. We define the *channel allocation* of a global type G to be the value of the function $\mathtt{ch}$.

Singleton allocation. The simplest channel allocation attributes a different channel to each communication occurring in the global type syntax tree. Formally, the singleton allocation is such that: $\forall \mathsf{n}, \mathsf{n}' \in G,\ \mathtt{ch}(\mathsf{n}) = \mathtt{ch}(\mathsf{n}') \iff \mathsf{n} = \mathsf{n}'$. Singleton allocations enjoy the following good properties.

Lemma 4.1 *For any global type G with singleton allocation, (1) G satisfies the linearity property; (2) for the finite channels k of G, $\mathscr{B}_k\langle G\rangle = 1$; (3) for the finite channels k of G, $\sum_k \mathscr{B}_k\langle G\rangle \leq size(G)$.*

Channel equalities. As well as values of the $\mathtt{ch}$ function, allocations can be seen as partitions of the set of nodes $\{\mathtt{n}\}_{\mathtt{n}\in G}$. We then define partition refinement through the notion of *channel equality*, i.e. the union of two partitions to produce a new allocation.

Definition 4.2 (channel equality) A *channel equality* is the substitution of two channels k and k' in a global type G by a single fresh channel k'' while keeping G linear.

As we take the singleton allocation as a base, we can describe channel allocations by sets E of channel equalities, the empty set corresponding to the singleton allocation. We write G_E the global type G with channel equalities E.

In the rest of this section, we always start from the singleton allocation and proceed by channel equality. We notably rely on the fact that the result of the equality of two finite channels is finite. Formally, if $\mathscr{B}_k\langle G\rangle = \infty$ then $\forall E, \mathscr{B}_k\langle G_E\rangle = \infty$.

Note that the total number of possible channel allocations is finite and corresponds to the number of partitions of a given finite set. The exact count (if we do not take into account the linearity property) is given by a Bell number [18] which is exponential in the size of the global type. Given the finite number of possible allocations, we know that there exists an algorithm to find allocations satisfying any decidable property. Notably, one can reach any given memory requirement (number of channels, buffer sizes).

Principal allocation. The most widely used allocation method attributes two communication channels (one in each direction) for each pair of participants. The session types in [1, 7] follow this allocation. Formally, the principal allocation is such that: $\forall \mathtt{n}, \mathtt{n}' \in G$ s.t. $\mathtt{pfx}(\mathtt{n}) = \mathtt{p} \to \mathtt{q} : k$ and $\mathtt{pfx}(\mathtt{n}') = \mathtt{p}' \to \mathtt{q}' : k', (k = k' \iff \mathtt{p} = \mathtt{p}' \wedge \mathtt{q} = \mathtt{q}')$.

Lemma 4.3 *For any global type G with principal allocation, (1) G satisfies the linearity property; (2) $chans(G) \leq n \times (n-1)$ where $n = prins(G)$.*

Greedy allocations. We now define a family of efficient algorithms, that give good allocation results in practice.

Definition 4.4 (Greedy allocation algorithm) *Given a global type with singleton allocation G, of initial node $\mathtt{n}_0$, and a successor function $\mathtt{succ}$ over the nodes, the function $\mathscr{I}^{\emptyset}_{\emptyset}(\mathtt{n}_0)$ is defined by:*

$$\mathscr{I}^{\mathrm{K}}_{E}(\mathtt{n}) = \mathscr{I}^{\mathrm{K}'}_{E'}(\mathtt{n}') \text{ where } \begin{cases} \mathtt{succ}(\mathtt{n}) = \mathtt{n}' \\ \mathtt{ch}(\mathtt{n}) = k \\ \mathrm{K}' = \mathrm{K} \cup \{k\} \end{cases} \wedge E' = \begin{cases} E \cup \{k = k'\} & \text{if } \exists k' \in \mathrm{K}, \mathrm{Lin}(G_{E\cup\{k=k'\}}) \\ E & \text{otherwise} \end{cases}$$
$$\mathscr{I}^{\mathrm{K}}_{E}(\mathsf{end}) = E$$

This algorithm is parameterised by the successor function over nodes (that can be given e.g. by a depth-first graph search) and by the choice between the possible channels $k' \in \mathrm{K}$ for equality. The greedy algorithm has the advantage of not backtracking and thus enjoys a polynomial complexity (if the choice procedures are polynomial) in the size of the graph. In particular, we define two efficient heuristics based on the generic greedy algorithm. In the greedy algorithm, we implement K by either:

1. (Early) a queue, so that we choose for channel equality the oldest channel $k' \in \mathrm{K}$.
2. (Late) a list, so that we choose for channel equality the latest channel $k' \in \mathrm{K}$.

The early and late allocations are not optimal in terms of total memory requirements (computed by $\sum_k \mathscr{B}_k\langle G\rangle$ when all channels are finite) but give good results in practice while being polynomial.

Example 4.5 (comparison of the allocations)
We apply the different allocation algorithms on a three-party stripped global type. The results are given in the adjacent table in term of number of allocated channels and total memory requirement. The greedy algorithms give the best results on this example, with the early greedy algorithm allocating less channels than the late greedy algorithm.

	Singleton	Principal	Early G.	Late G.
$\mathsf{n}_0 : \mathtt{A} \to \mathtt{B};$	k_0	k_0	k_0	k_0
$\mathsf{n}_1 : \mathtt{B} \to \mathtt{A};$	k_1	k_1	k_1	k_1
$\mathsf{n}_2 : \mathtt{A} \to \mathtt{B};$	k_2	k_0	k_0	k_0
$\mathsf{n}_3 : \mathtt{A} \to \mathtt{B};$	k_3	k_0	k_0	k_1
$\mathsf{n}_4 : \mathtt{A} \to \mathtt{C};$	k_4	k_2	k_2	k_2
$\mathsf{n}_5 : \mathtt{C} \to \mathtt{B};$	k_5	k_3	k_1	k_3
$\mathsf{n}_6 : \mathtt{B} \to \mathtt{C};$	k_6	k_4	k_2	k_1
$\mathsf{n}_7 : \mathtt{B} \to \mathtt{C}$	k_7	k_4	k_2	k_2
Nb channels	8	5	3	4
Memory Req.	8	7	5	5

5 Global and Local Refinements

5.1 Global Refinement: Insertion of Confirmation Messages

This subsection presents two algorithms to automatically insert confirmation messages to limit the buffer size requirements of a given session type. This allows the treatment of infinite channels left out in the previous section.

Our algorithms work in three steps: first, (1) the confirmation messages to introduce are computed from the global type; next, (2) the algorithms determine how many unfoldings of the global type are necessary to reach the desired buffer sizes; finally, (3) the confirmation messages are inserted in the unfolded global type.

For step (1), *the optimal-confirmation algorithm* looks for the minimal number of confirmation messages to introduce to bound all channel buffers. The algorithm will in practice try every combination of sender, receiver and channel. Note that at most one message per infinite channel is needed. *The instant-confirmation algorithm* relies on this property and adds a confirmation for each infinite channel of the original session type. For a given channel k whose latest prefix is $\mathtt{p} \to \mathtt{p}' : k$, the confirmation is of the form $\mathtt{p}' \to \mathtt{p} : k'$ with k' chosen by one of the channel allocation algorithms.

For step (2), once the set of confirmation messages is chosen, both algorithms determine the optimal number of unfolding of the global type needed to reach the desired buffer sizes (the unfolding happens within the recursion calls: $\phi^1(\mu\mathbf{x}.G) = \mu\mathbf{x}.\phi^1(G[\mu\mathbf{x}^*.G/\mathbf{x}])$, $\phi^1(\mu\mathbf{x}^*.G) = G$ and $\phi^{n+1}(G) = \phi^n(\phi^1(G))$). After each unfolding and until the memory targets are met, the confirmation messages are introduced (at the level of recursion variables for the instant version) and buffer sizes are checked.

Example 5.1 Recall $G = \mu\mathbf{x}.\mathsf{Alice} \to \mathsf{Bob}\colon s_1\,\langle\mathsf{nat}\rangle;\mathsf{Carol} \to \mathsf{Alice}\colon s_3\,\langle\mathsf{bool}\rangle;\mathbf{x}$ (Example (b) from § 1) where the channels s_1 and s_3 are infinite. We illustrate the optimal confirmation algorithm on the l.h.s. and the instant one on the r.h.s for G below. The two algorithms give the following results when asked to limit buffer sizes to 2.

$$
\begin{array}{rl}
\mu\mathbf{x}. & \mathsf{Alice} \to \mathsf{Bob}\colon s_1\,\langle\mathsf{nat}\rangle; \\
& \mathsf{Carol} \to \mathsf{Alice}\colon s_3\,\langle\mathsf{bool}\rangle; \\
& \mathsf{Alice} \to \mathsf{Bob}\colon s_1\,\langle\mathsf{nat}\rangle; \\
\mathsf{n}_0 : & \mathsf{Bob} \to \mathsf{Carol}\colon s_4\,\langle\mathsf{unit}\rangle; \\
& \mathsf{Carol} \to \mathsf{Alice}\colon s_3\,\langle\mathsf{bool}\rangle;\mathbf{x}
\end{array}
\qquad
\begin{array}{rl}
\mu\mathbf{x}. & \mathsf{Alice} \to \mathsf{Bob}\colon s_1\,\langle\mathsf{nat}\rangle; \\
& \mathsf{Carol} \to \mathsf{Alice}\colon s_3\,\langle\mathsf{bool}\rangle; \\
& \mathsf{Alice} \to \mathsf{Bob}\colon s_1\,\langle\mathsf{nat}\rangle; \\
& \mathsf{Carol} \to \mathsf{Alice}\colon s_3\,\langle\mathsf{bool}\rangle; \\
\mathsf{n}_1 : & \mathsf{Bob} \to \mathsf{Alice}\colon s_4\,\langle\mathsf{unit}\rangle; \\
\mathsf{n}_2 : & \mathsf{Alice} \to \mathsf{Carol}\colon s_5\,\langle\mathsf{unit}\rangle;\mathbf{x}
\end{array}
$$

The *optimal-confirmation* discovers (on the l.h.s.) that n_0 is enough to bound all channels. On the r.h.s., the *instant-confirmation* introduces two confirmation messages, n_1 and n_2. In both, unfolding is done exactly once to reach a bound of 2 on all buffer sizes.

5.2 Local Refinement: Messaging Optimisations

One of the significant practical concerns in systems with messaging is to optimise interactions through more asynchronous data processing to increase parallelism. In § 5.1, we study how to limit the size of buffers by refining global types. Our next step concerns messaging optimisations that respect the agreed buffer sizes. Our recent work [13, 14] developed a new form of subtyping, the *asynchronous subtyping*, that characterises the compatibility between classes of type-safe permutations of actions, in order to send messages before receiving. This subtyping allows, however, not only Bob_5 in (1.1) in § 1 but also $\mu X.s_2!\langle Orange\rangle;X$ as a refinement of Bob, which changes all buffer sizes from 1 to ∞, leading to buffer overflows. Our aim is to overcome this problem by controlling permutations *locally* with the help of the IO-dependency analysis. The key idea is to prohibit the permutation of an output action at k_0 with an input or branching action which prevents (by IO-causality) the accumulation of values in k_0.

Recall Definition 3.4. We define the *minimal resetting paths* to be the paths that satisfy the reset property while none of their suffix does. Then, we define the *dependent nodes* of channel k, noted $\mathtt{dep}(k)$ to be the union of the reset nodes of the minimal resetting paths that end with k. This set of nodes characterises a buffer usage.

First, for a given G, we choose a partition $\{N_0,\ldots,N_n\}$ of the set of nodes of G. This partition should satisfy the two properties: $\forall \mathsf{n}\in N_i, \mathsf{n}'\in N_j, \mathtt{ch}(\mathsf{n})=\mathtt{ch}(\mathsf{n}')\Rightarrow N_i=N_j$ and $\forall \mathsf{n}\in N_i, \mathtt{dep}(\mathtt{ch}(\mathsf{n}))\subset N_i$. The choice of a partitioning depends in particular on a choice of reset and dependent nodes. Note that the trivial partitioning (with only one partition) is always possible. Since that, for each channel k, all nodes using k are part of the same partition (written $N(k)$), we can annotate all uses of k in G by $N(k)$.

In the example below, the partitioning is made of $N_1=\{\mathsf{n}_1,\mathsf{n}_2\}$ and $N_2=\{\mathsf{n}_3,\mathsf{n}_4\}$: we give the annotated session graph (with the IO-dependencies highlighted) on the left and the projected types (where the annotations are kept) on the right.

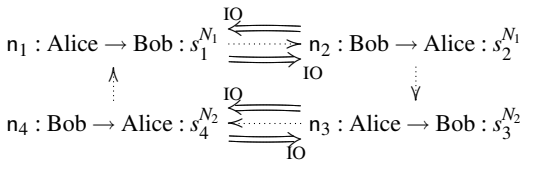

$$T_{\text{Alice}}=\mu\mathbf{x}.s_1^{N_1}!;s_2^{N_1}?;s_3^{N_2}!;s_4^{N_2}?;\mathbf{x}$$
$$T_{\text{Bob}}=\mu\mathbf{x}.s_1^{N_1}?;s_2^{N_1}!;s_3^{N_2}?;s_4^{N_2}!;\mathbf{x}$$
$$T_{\text{Alice}}^{opt}=\mu\mathbf{x}.s_1^{N_1}!;s_3^{N_2}!;s_2^{N_1}?;s_4^{N_2}?;\mathbf{x}$$

Next, we apply the size-preserving asynchronous communication subtyping, following the annotations on the projected types. The relation $T\ll T'$ means T is more asynchronous than (or more optimised than) T'. The main rule is:

$$(\mathsf{OI})\qquad k^N!\langle U\rangle;k_0^{N_0}?\langle U'\rangle;T\ \ll\ k_0^{N_0}?\langle U'\rangle;k^N!\langle U\rangle;T\qquad (N\cap N_0=\emptyset)$$

where the two prefixes are permutable if the IO-chains of the two prefixes are disjoint. We can *always* permute two inputs and two outputs with distinct channels since they do not contribute to the input and output alternations that constitute the IO-chains.

The branching/selection rules are similarly defined, and others are context rules. Then we define a coinductive subtyping relation $T_1 \leqslant_c T_2$ as a form of type simulation, following [13, 14]. The important fact is that $\leqslant_c$ does not alter buffer sizes: suppose $[\![G]\!] = \{T@\mathtt{p}\}_{\mathtt{p}}$ with $T@\mathtt{p} = (G \upharpoonright \mathtt{p})@\mathtt{p}$ and $\mathtt{p} \in G$. Assume $T@\mathtt{p} \leqslant_c T'@\mathtt{p}$ with $[\![G']\!] = \{T'@\mathtt{p}\}_{\mathtt{p}}$ Then $\mathscr{B}_k\langle G\rangle = \mathscr{B}_k\langle G'\rangle$. Since there is no change in the buffer bounds, Type and Buffer Safety are just proved from Theorem 3.8 and Corollary 3.9.

In the example above, in Alice's type, we can permute $s_2^{N_1}?$ and $s_3^{N_2}!$ ($T^{opt}_{\mathrm{Alice}} \leqslant_c T_{\mathrm{Alice}}$) since $N_1 \cap N_2 = \emptyset$, keeping the size of each buffer one. Hence process typable by T^{opt}_{Alice} can safely send message at s_3 before input at s_2. In Alice-Bob$_5$-Carol from § 1, the original global type G annotated by IO-chains has only one partition $N = \{\mathsf{n}_1, \mathsf{n}_2, \mathsf{n}_3\}$:

$$\mu\mathbf{x}.\ \mathrm{Alice} \to \mathrm{Bob}\colon s_1^N\,\langle\mathsf{nat}\rangle; \mathrm{Bob} \to \mathrm{Carol}\colon s_2^N\,\langle\mathsf{string}\rangle; \mathrm{Carol} \to \mathrm{Alice}\colon s_3^N\,\langle\mathsf{real}\rangle; \mathbf{x}$$

Bob's local type is $\mu\mathbf{x}.s_1^N?\langle\mathsf{nat}\rangle; s_2^N!\langle\mathsf{string}\rangle; \mathbf{x}$, which prevents any optimisation $\ll$ by (OI). Hence, Bob$_5$ is not typable. Some typable examples are given in the next section.

6 Application Examples

Multi-buffering algorithm. The double buffering algorithm [5] is widely used in high-performance and multicore computing. We generalise this algorithm to *multi-buffering* [15], and solve an open issue in our previous work [14, § 5]. The aim is to transport a large amount of data as a series of units (say each unit is 16kB) from a source (Source) to a transformer (called Kernel). Each unit gets processed at Kernel and delivered to a sink (Sink). Kernel uses n 16kB buffers, named B_i, to maximise the message transfer asynchrony. Processes which represent Source, Sink, Kernel and Optimised Kernel are given below using parameterised processes [21] (i.e. where $\mathtt{foreach}(i<n)\{P[i]\}$ means we iterate $P[i]$ for $0 \leq i < n$):

Source: $\mu X.\mathtt{foreach}(i<n)\{r_i?();\ s_i!\langle y_i\rangle\};X$ *Sink:* $\mu X.\mathtt{foreach}(i<n)\{t_i!\langle\rangle; u_i?(z_i)\};X$

Kernel: $\mu X.\mathtt{foreach}(i<n)\{r_i!\langle\rangle; s_i?(x_i); t_i?(); u_i!\langle x_i\rangle\};X$

Optimised Kernel: $r_0!\langle\rangle;\ldots;r_{n-1}!\langle\rangle;\mu X.\mathtt{foreach}(i<n)\{s_i?(x_i); t_i?(); u_i!\langle x_i\rangle; r_i!\langle\rangle\};X$

In the loop, Kernel notifies Source with signals at r_i that it is ready to receive data in each channel s_i of buffer B_i. Source complies, sending one unit via s_i. Then Kernel waits for Sink to inform (via t_i) that Sink is ready to receive data via u_i: upon receiving the signals, Kernel sends the unit of processed data to Sink via u_i. If Kernel sends the n notifications to $r_0,\ldots,r_{n-1}$ *ahead* like Optimised Kernel, Source can start its work for the next unit (sending y_j at s_j) without waiting for other buffers.

The following proposition means that *the n-buffers of size one in Kernel simulate one buffer of size n*, maximising the asynchrony. The proof is done by annotating the global type with partitions $\{r_i, s_i\}$ and $\{u_i, t_i\}$, and checking that the permutation of the projected Kernel type satisfies the (OI) rule.

Proposition 6.1 (n-buffering correctness) *Source-Optimal Kernel-Sink satisfies both progress and communication-safety. Also each buffer at s_i and u_i holds at most one unit.*

If Optimised Kernel is optimised as $r_0!\langle\rangle;\ldots;\texttt{foreach}(i<n)\{r_i!\langle\rangle;s_i?(x_i);t_i?;u_i!\langle x_i\rangle\}$ (which is not typable in our system), then all buffers are forced to hold 2 *units*. This unsafe optimisation is typable in [14] but prevented here by Proposition 6.1.

MPSoC buffer allocations. We take a use case from [4] and present performance optimisations of MPSoC. We show that this use case can be described with branching and parameterised types (from [21]) and can then be optimised and verified by applying all of the methods developed in § 3–5.2.

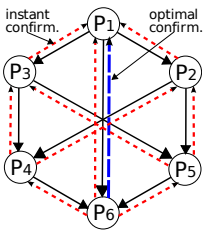

The diagram shows six processes connected by nine FIFOs and its specification is given below:

$$G = \mu\mathbf{x}.(\ \texttt{foreach}(i\in\{2,3\})\{G[i]\};\mathtt{p}[1]\to\mathtt{p}[6]:\mathsf{on});\mathbf{x}$$

$$G[i] = \mathtt{p}[1]\to\mathtt{p}[i]:\left\{\begin{array}{l}\mathsf{on}:\mathtt{p}[i]\to\mathtt{p}[4]:\mathsf{on};\ \mathtt{p}[4]\to\mathtt{p}[6]:\mathsf{on};\\ \mathsf{off}:\mathtt{p}[i]\to\mathtt{p}[5]:\mathsf{off};\ \mathtt{p}[5]\to\mathtt{p}[6]:\mathsf{off}\end{array}\right\}$$

where $\texttt{foreach}(i\in I)\{G[i]\}$ means that $G[i]$ is repeated for each $i\in I$.

For simplicity, we perform the principal allocation (§ 4). By Definition 3.2, buffers of infinite sizes are needed if the data is continuously sent by $\mathtt{p}[1]$. To limit their upper bounds to size m, we apply the two global refinement algorithms presented in § 5.1. The optimal algorithm inserts one confirmation after each m iterations as in: $\mu\mathbf{x}.\texttt{foreach}(n<m)\{\texttt{foreach}(i\in\{2,3\})\{G[i]\}\};\mathtt{p}[6]\to\mathtt{p}[1]\langle\mathsf{unit}\rangle:\mathbf{x}$. The instant confirmation inserts nine messages (as shown by the red arrows), but has the advantage of allowing more asynchronous optimisation (§ 5.2). For example, using instant confirmation, $\mathtt{p}[1]$ can send the $m+1$-th message to $\mathtt{p}[2]$ without waiting for the confirmation of the m-th message from $\mathtt{p}[6]$, while this is impossible with optimal confirmation. Finally we can prove that the above MPSoC six processes are communication-safe, and satisfy progress and buffer-safety (after global and local refinements).

7 Related Work

Checking buffer bounds based on global specifications has been studied through Petri nets and Synchronous data flow. Recent advances [8] in the study of Kahn Process Networks (KPN) have improved Parks's algorithm [17] to ensure safe executions of stream-based applications with bounded buffers, using an appropriate scheduling policy. Their theory is applied to KPN applications on MPSoC [4], demonstrating the effectiveness of non-uniform, fine-grained buffer allocations. By contrast, our approach is type-based and relies on the existence of a global specification that brings additional guarantees (such as deadlock-freedom) and allows global choreography manipulation and refinements. It is moreover directly applicable to programming languages [11, 21] by extending existing type syntax and checking.

The idea of using a type-abstraction to investigate channel communications goes back to Nielson & Nielson's work on CML [16]. Gay & Vasconcelos [7] propose a linear type system for binary sessions to enforce buffer bounds computed by a fixed point method. Their work is thus limited to a particular channel allocation (i.e. principal, cf. § 4) and does not extend to multiparty interactions (their method would find that the

buffers in Example (a) are infinite). Terauchi & Megacz [20] describe a polynomial method to infer buffer bounds of a concurrent language through program analysis using linear programming techniques, improving on previous work in [12], see [20, § 7]. Our bound computation method differs in that it starts from a direct type-based abstraction of global interaction structures, namely session graphs, not from direct investigation of local types nor processes (normally in distributed systems, a peer does not know other peer's type or implementation [11]). It also leads to the general simplicity of the analysis, and the uniform treatment of subtle issues such as asynchronous optimisations. Thanks to session types, the channel passing problem in [20, § 6] does not arise in our analysis: different (possibly newly generated) sessions and names can be stored in the same buffer, still giving the exact bound of stored channels. None of [7, 20] have studied either channel allocation, global refinement or messaging optimisation.

Among process calculi for service-oriented computing (SOC), contracts [3] and the conversation calculus [2] provide static type checking for a series of interactions and ensure progress. We demonstrate the advantage of global types by the simplicity of our analysis and the uniform treatments and articulation of our various algorithms. Our approach is, however, extensible to these calculi because (1) the IO-causality analysis does not rely on the form of session branches so that other form of sums can be analysed by the same technique; and (2) combining with a polynomial inference which builds a graph from a collection of local types $[\![G]\!]$ [14], Subject Reduction Theorem can be proved using our invariance method noting that we use $[\![G]\!]$ for the proofs. An extension to other formalisms for SOC including [2, 3] is an interesting future work.

Further topics include the enrichment of global types with more quantitative information (such as distance, probabilities and weights), which would enable finer-grained analyses and optimisations.

References

1. L. Bettini, M. Coppo, L. D'Antoni, M. D. Luca, M. Dezani-Ciancaglini, and N. Yoshida. Global progress in dynamically interleaved multiparty sessions. In *CONCUR*, volume 5201 of *LNCS*, pages 418–433, 2008.
2. L. Caires and H. T. Vieira. Conversation types. In *ESOP*, volume 5502 of *LNCS*, pages 285–300. Springer, 2009.
3. G. Castagna and L. Padovani. Contracts for mobile processes. In *CONCUR*, number 5710 in LNCS, pages 211–228, 2009.
4. E. Cheung, H. Hsieh, and F. Balarin. Automatic buffer sizing for rate-constrained KPN applications on multiprocessor system-on-chip. In *HLDVT'07*, pages 37–44. IEEE, 2007.
5. A. Donaldson, D. Kroening, and P. Rmmer. Automatic analysis of scratch-pad memory code for heterogeneous multicore processors. In *TACAS*, LNCS 6015, pages 280–295, 2010.
6. M. Fähndrich et al. Language support for fast and reliable message-based communication in singularity OS. In *EuroSys2006*, ACM SIGOPS, pages 177–190. ACM Press, 2006.
7. S. Gay and V. T. Vasconcelos. Linear type theory for asynchronous session types. *JFP*, 2009.
8. M. Geilen and T. Basten. Requirements on the Execution of Kahn Process Networks. In *ESOP*, volume 2618 of *LNCS*, pages 319–334. Springer, 2003.
9. K. Honda, V. T. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, volume 1381 of *LNCS*, pages 22–138. Springer, 1998.

10. K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL*, pages 273–284, 2008.
11. R. Hu, N. Yoshida, and K. Honda. Session-Based Distributed Programming in Java. In *ECOOP'08*, volume 5142 of *LNCS*, pages 516–541, 2008.
12. N. Kobayashi, M. Nakade, and A. Yonezawa. Static analysis of communication for asynchronous concurrent programming languages. In *SAS*, LNCS 983, pages 225–242, 1995.
13. D. Mostrous and N. Yoshida. Session-Based Communication Optimisation for Higher-Order Mobile Processes. In *TLCA*, volume 5608 of *LNCS*, pages 203–218. Springer, 2009.
14. D. Mostrous, N. Yoshida, and K. Honda. Global principal typing in partially commutative asynchronous sessions. In *ESOP'09*, volume 5502 of *LNCS*, pages 316–332, 2009.
15. Multi-buffering. `http://en.wikipedia.org/wiki/Multiple_buffering`.
16. H. Nielson and F. Nielson. Higher-order concurrent programs with finite communication topology. In *POPL*, pages 84–97, 1994.
17. T. Parks. *Bounded Scheduling of Process Networks*. PhD thesis, California Barkeley, 1995.
18. G.-C. Rota. The number of partitions of a set. *Amer. Math. Monthly*, 71:498–504, 1964.
19. K. Takeuchi, K. Honda, and M. Kubo. An interaction-based language and its typing system. In *PARLE'94*, volume 817 of *LNCS*, pages 398–413. Springer, 1994.
20. T. Terauchi and A. Megacz. Inferring channel buffer bounds via linear programming. In *ESOP*, volume 4960 of *LNCS*, pages 284–298. Springer, 2008.
21. N. Yoshida, P.-M. Deniélou, A. Bejleri, and R. Hu. Parameterised multiparty session types. In *FoSSaCs*, volume 6014 of *LNCS*, pages 128–145, 2010.

A Appendix for Section 2: Reductions and Structural Congruence

We define the full set of reduction and structural congruence rules mentioned in Section 2.

$$
\begin{array}{rlll}
\overline{a}_{[2..n]}(\tilde{s}^{\tilde{n}}).P_1 \mid a_{[2]}(\tilde{s}).P_2 \mid \cdots \mid a_{[n]}(\tilde{s}).P_n & & & \\
\quad \to (\nu\,\tilde{s})(P_1 \mid P_2 \mid \ldots \mid P_n \mid s_1^{n_1}:\emptyset \mid \ldots \mid s_m^{n_m}:\emptyset) & & & [\textsc{Link}] \\
s!\langle \tilde{e} \rangle; P \mid s^n : \tilde{h} & \to P \mid s^n : \tilde{h} \cdot \tilde{v} & (n \gtrsim |\tilde{h}|,\ e_i \downarrow v_i) & [\textsc{Send}] \\
s!\langle\!\langle \tilde{t} \rangle\!\rangle; P \mid s^n : \tilde{h} & \to P \mid s^n : \tilde{h} \cdot \tilde{t} & (n \gtrsim |\tilde{h}|) & [\textsc{Deleg}] \\
s \lhd l; P \mid s^n : \tilde{h} & \to P \mid s^n : \tilde{h} \cdot l & (n \gtrsim |\tilde{h}|) & [\textsc{Sel}] \\
s?(\tilde{x}); P \mid s^n : \tilde{v} \cdot \tilde{h} & \to P[\tilde{v}/\tilde{x}] \mid s^n : \tilde{h} & & [\textsc{Recv}] \\
s?(\!(\tilde{t})\!); P \mid s^n : \tilde{t} \cdot \tilde{h} & \to P \mid s^n : \tilde{h} & & [\textsc{SRec}] \\
s \rhd \{l_i : P_i\}_{i \in I} \mid s^n : l_j \cdot \tilde{h} & \to P_j \mid s^n : \tilde{h} & (j \in I) & [\textsc{Bra}] \\
\mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q & \to P & (e \downarrow \mathsf{true}) & [\textsc{IfT}] \\
\mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q & \to Q & (e \downarrow \mathsf{false}) & [\textsc{IfF}] \\
P \to P' & \Rightarrow (\nu\, n)P \to (\nu\, n)P' & & [\textsc{Scop}] \\
P \to P' & \Rightarrow P \mid Q \to P' \mid Q & & [\textsc{Par}] \\
P \equiv P' \text{ and } P' \to Q' \text{ and } Q' \equiv Q & \Rightarrow P \to Q & & [\textsc{Str}]
\end{array}
$$

where $n = a$ or $n = \tilde{s}$.

$$
\begin{array}{c}
P \mid \mathbf{0} \equiv P \qquad P \mid Q \equiv Q \mid P \qquad (P \mid Q) \mid R \equiv P \mid (Q \mid R) \\
(\nu\, n)P \mid Q \equiv (\nu\, n)(P \mid Q) \ \text{ if } n \notin \mathsf{fn}(Q) \quad (\nu\, nn')P \equiv (\nu\, n'n)P \\
(\nu\, n)\mathbf{0} \equiv \mathbf{0} \quad \mu X.P \equiv P[\mu X.P/X] \quad (\nu\, s_1..s_n)\Pi_i s_i^{n_i} : \emptyset \equiv \mathbf{0}
\end{array}
$$

n denotes either a or s. For process P, $\mathsf{fpv}(P)$ and $\mathsf{fn}(P)$ respectively denote the sets of *free process variables* and *free identifiers* in P. A sequence of parallel composition is written $\Pi_i P_i$.

B Appendix for Sections 3.1 and 3.2: Bound Analysis in Multiparty Sessions

This section lists the omitted definitions of Sections 3.1 and 3.2 and details causal dependencies.

B.1 Dependency Analysis Examples

We list all three dependencies together for the reader's convenience.

Definition B.1 (dependencies) *Fix G. The relations $\prec_\phi$, with $\phi \in \{\mathsf{II}, \mathsf{IO}, \mathsf{OO}\}$, over ordered prefixes are defined by:*

$$\mathsf{n}_1 \prec_{\mathtt{II}} \mathsf{n}_2 \text{ if } \mathsf{n}_1 \prec \mathsf{n}_2 \wedge \begin{cases} \mathsf{n}_1 = \mathtt{p}_1 \to \mathtt{p} : k_1 \\ \mathsf{n}_2 = \mathtt{p}_2 \to \mathtt{p} : k_2 \end{cases} \text{ s.t. } \mathtt{p}_1 \neq \mathtt{p}_2 \Leftrightarrow k_1 \neq k_2$$

$$\mathsf{n}_1 \prec_{\mathtt{IO}} \mathsf{n}_2 \text{ if } \mathsf{n}_1 \prec \mathsf{n}_2 \wedge \begin{cases} \mathsf{n}_1 = \mathtt{p}_1 \to \mathtt{p} : k_1 \\ \mathsf{n}_2 = \mathtt{p} \to \mathtt{p}_2 : k_2 \end{cases} \text{ with } k_1 \neq k_2$$

$$\mathsf{n}_1 \prec_{\mathtt{OO}} \mathsf{n}_2 \text{ if } \mathsf{n}_1 \prec \mathsf{n}_2 \wedge \begin{cases} \mathsf{n}_1 = \mathtt{p} \to \mathtt{p}_1 : k_1 \\ \mathsf{n}_2 = \mathtt{p} \to \mathtt{p}_2 : k_2 \end{cases} \text{ s.t. } \mathtt{p}_1 \neq \mathtt{p}_2 \Rightarrow k_1 \neq k_2$$

The relation $\prec_{\mathtt{II}}$ expresses the fact that a principal follows the receive order when different senders send on different channels, or when a unique sender sends two messages on the same channel. The relation $\prec_{\mathtt{IO}}$ asserts the order between a reception by a principal and the next message it sends. Finally, the relation $\prec_{\mathtt{OO}}$ orders consecutive sends by a same principal, except when a unique channel is used for different concurrent receivers. We now combine these relations to define *dependency relations* that express causality chains.

Definition B.2 (dependency relations)

1. *An* input dependency (I-dependency) from n_1 to n_2 *is a chain* $\mathsf{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathsf{n}_2$ *($n \geq 0$) such that $\phi_i = \mathsf{I0}$ for $1 \leq i \leq n-1$ and $\phi_n = \mathsf{II}$.*
2. *An* output dependency (O-dependency) from n_1 to n_2 *is a chain* $\mathsf{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathsf{n}_2$ *($n \geq 1$) such that $\phi_i \in \{\mathsf{OO}, \mathsf{IO}\}$.*
3. *An* input-output dependency (IO-dependency) from n_1 to n_2 *is a chain* $\mathsf{n}_1 \prec_{\phi_1} \cdots \prec_{\phi_n} \mathsf{n}_2$ *($n \geq 1$) such that $\phi_i = \mathsf{IO}$.*

The O-dependency characterises the sequentiality between exchanged messages while the input dependency makes sure that a reception at n_2 is guarded and does not conflict with the reception at n_1. In (2), the last II-ordering is needed since, if it ends with an IO-edge, an input at n_2 may not have been completed.

Example B.3 We illustrate the notion of infinite channels on a variation from our running example.

$$\mu \mathbf{x}.\text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \text{Carol} \to \text{Alice}: s_3 \text{Alice} \to \text{Carol}: s_2\,\langle \mathsf{string} \rangle; \mathbf{x}$$

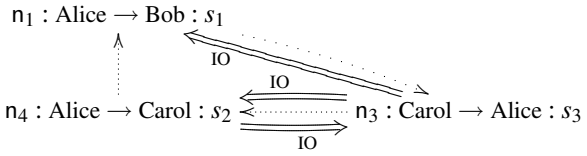

Since the node n_1 is part of the main cycle $\mathsf{n}_1\mathsf{n}_3\mathsf{n}_4$, but not part of the IO-cycle $\mathsf{n}_3\mathsf{n}_4$, the channel s_1 is infinite.

Proposition B.4 *Computing all IO-dependencies of a given session graph has complexity $O(size(G)^3)$.*

Proof. We explore the graph up to one unfolding. For each explored node, we find the IO-dependencies that end there. The overall complexity is thus $O(size(G)^3)$.

Example B.5 (reset function) In the following session type (part of the previous one-unfolding of the session involving Alice, Bob and Carol), we are interested in the truth value of Reset($\tilde{\mathsf{n}}$,) where $\tilde{\mathsf{n}} = \mathsf{n}_1 \prec \mathsf{n}_2 \prec \mathsf{n}_3 \prec \mathsf{n}_4$:

$$\begin{array}{ll} \mathsf{n}_1: & \text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \\ \mathsf{n}_2: & \text{Bob} \to \text{Carol}: s_2\,\langle \mathsf{string} \rangle; \\ \mathsf{n}_3: & \text{Carol} \to \text{Alice}: s_3\,\langle \mathsf{real} \rangle; \\ \mathsf{n}_4: & \text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \mathsf{end} \end{array}$$

The only node of $\tilde{\mathsf{n}}$ using the same channel as n_4 is n_1. Since $\mathsf{n}_1 \prec_{\mathtt{IO}} \mathsf{n}_2 \prec_{\mathtt{IO}} \mathsf{n}_3 \prec_{\mathtt{IO}} \mathsf{n}_4$, there is an IO-dependency between n_4 and n_1. The property Reset$(\tilde{\mathsf{n}}, \mathsf{n}_4)$ holds.

This example illustrates that Bob has to complete the reception on channel s_1 in n_1 since he is the sender of the next message. The same goes for Carol and Alice, thus preventing Bob from accumulating in s_1 the two messages sent in n_1 and n_4. We therefore know that the buffer of channel s_1 is empty prior to the message of n_4.

Note finally that in the following session type, the property Reset$(\mathsf{n}_1\mathsf{n}_2\mathsf{n}_3\mathsf{n}_4\mathsf{n}_5,)$ also holds, although the buffer of s_1 can accumulate a second value in n_3.

$$\begin{array}{ll} \mathsf{n}_1: & \text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \\ \mathsf{n}_2: & \text{Bob} \to \text{Carol}: s_2\,\langle \mathsf{string} \rangle; \\ \mathsf{n}_3: & \text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \\ \mathsf{n}_4: & \text{Carol} \to \text{Alice}: s_3\,\langle \mathsf{real} \rangle; \\ \mathsf{n}_5: & \text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \mathsf{end} \end{array}$$

Proposition B.6 *Computing the buffer bounds of a given global session graph is polynomial.*

Proof. The buffer bound computation algorithm is polynomial because the lengths of the paths it has to examine are limited. The reason is that, for a given linear session

graph, any IO-chain (such as the ones mentionned in the reset property definition) that is longer than twice the number of participants imply the existence of a shorter IO-sub-chain.

As a matter of fact, if a given participant appears twice as a receiver and once as a sender in a minimal IO-chain, it means that one of the two receptions use the same channel as the send done by this participant (by definition of the IO-dependency). Any other situation (three receptions) implies that the IO-chain can be short-cut.

As the search for an IO-chains can be limited to chains that are no longer than twice the number of participants, paths that need to be explored do not need to be longer than $O(n * size(G))$ (with n the number of participants). Consequently, the paths that are explored by the bound computation algorithm are guaranteed to satisfy the reset property when they reach a certain length that depends only linearly on the size of the graph.

This ensures that the in-depth exploration of the session graph that is done to compute the buffer bounds is polynomial in the size of the graph (assuming the number of participants is a constant).

C Appendix for Section 3.3: Typing Rules and Buffer Safety Theorem

This section lists the typing system omitted from the main section and gives the proofs of the subject reduction and Buffer safety.

We start from the example of the linearity.

Example C.1 (linearity) We illustrate the linearity property on a counter-example.

$$\begin{array}{ll} \mathsf{n}_1: & \text{Alice} \to \text{Bob}: s_1\,\langle \mathsf{nat} \rangle; \\ \mathsf{n}_2: & \text{Carol} \to \text{Bob}: s_1\,\langle \mathsf{string} \rangle; \end{array}$$

In the above excerpt of a global type, linearity is broken since there is no input and output causality between the two prefixes which have the same channel s_1. Actually, Bob faces a race condition since both Alice and Carol try to enqueue the values on the same channel. On the other hand, the previous global type in Example B.5 satisfies the linearity condition since there is the O-dependency from n_1 to n_4 such that $\mathsf{n}_1 \prec_{\mathtt{IO}} \mathsf{n}_2 \prec_{\mathtt{IO}} \mathsf{n}_3 \prec_{\mathtt{IO}} \mathsf{n}_4$ and an I-dependency $\mathsf{n}_1 \prec_{\mathtt{II}} \mathsf{n}_4$. With these I and O chains, we can ensure no race at s_1. More examples about linearity can be found in [10].

C.1 Projection

The following defines the projection of a global type to local types at each participant. Then the *projection of G onto* $\mathtt{p}$, written $G \upharpoonright \mathtt{p}$, is inductively given as:

- $(\mathtt{p}_1 \to \mathtt{p}_2 : k\,\langle U \rangle; G') \upharpoonright \mathtt{p} =$
$$\begin{cases} k!\langle U \rangle.(G' \upharpoonright \mathtt{p}) & \text{if } \mathtt{p} = \mathtt{p}_1 \neq \mathtt{p}_2 \\ k?\langle U \rangle.(G' \upharpoonright \mathtt{p}) & \text{if } \mathtt{p} = \mathtt{p}_2 \neq \mathtt{p}_1 \\ (G' \upharpoonright \mathtt{p}) & \text{if } \mathtt{p} \neq \mathtt{p}_2 \text{ and } \mathtt{p} \neq \mathtt{p}_1 \end{cases}$$

– $(\mathtt{p}_1 \to \mathtt{p}_2 \colon k\{l_j \colon G_j\}_{j\in J}) \upharpoonright \mathtt{p} =$
$$\begin{cases} \oplus\{l_j \colon (G_j \upharpoonright \mathtt{p})\}_{j\in J} & \text{if } \mathtt{p} = \mathtt{p}_1 \neq \mathtt{p}_2 \\ \&\{l_j \colon (G_j \upharpoonright \mathtt{p})\}_{j\in J} & \text{if } \mathtt{p} = \mathtt{p}_2 \neq \mathtt{p}_1 \\ (G_1 \upharpoonright \mathtt{p}) & \text{if } \mathtt{p} \neq \mathtt{p}_2 \text{ and } \mathtt{p} \neq \mathtt{p}_1 \\ & \text{and } \forall i,j \in I. G_i \upharpoonright \mathtt{p} = G_j \upharpoonright \mathtt{p} \end{cases}$$
– $(\mu \mathbf{x}.G) \upharpoonright \mathtt{p} = \mu \mathbf{x}.(G \upharpoonright \mathtt{p})$, $\mathbf{x} \upharpoonright \mathtt{p} = \mathbf{x}$, and $\mathsf{end} \upharpoonright \mathtt{p} = \mathsf{end}$.

When a side condition does not hold the map is undefined. We also use the extended projection from [21] for branching (for which the proof does not change).

C.2 Typing System for Initial Processes.

We list the typing system for initial processes in figure 1.

C.3 Typing rules for runtime processes

Typing runtime processes is more involved since, for subject reduction, we need to maintain an invariant about the number of messages in buffers and processes. We start by extending the types and environments to accommodate buffer sizes:

$$\Delta \ ::= \ \emptyset \mid \Delta, \tilde{s}^{\tilde{n}} \colon \{H_{\mathtt{p}}@\mathtt{p}\}_{\mathtt{p}\in I} \qquad M \ ::= \ k!\langle U\rangle \mid k \oplus l \mid M;M \qquad H \ ::= \ T \mid M \mid M;T$$

where $M, M', \ldots$ range over *message types* which represent the messages contained in buffers; $H, H', \ldots$ range over *generalised types* which are either local types, message types, or message types followed by local types.

We first type a single queue (buffer), in which the turnstile $\vdash$ is decorated with s (where s is the session of the current buffer) and where the session environments are mappings from channels to message types. We list the full typing rules for buffers in figure 2.

Since all rules are similar to standard multiparty session typing systems, we only explain two rules from figure 2.

$$\frac{}{\Gamma \vdash_{\{s_k\}} s_k^{\tilde{n}} : \emptyset \triangleright \emptyset} \qquad \frac{\Gamma \vdash v_i : S_i \qquad \Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \triangleright \Delta}{\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \cdot \tilde{v} \triangleright \Delta; \tilde{s}^{\tilde{n}} : k!\langle \tilde{S}\rangle @ \mathtt{p}}$$

The empty buffer has an empty session environment. Each message adds an output type to the current channel type where ; is defined by:

$$\Delta; \tilde{s}^{\tilde{m}} : M@\mathtt{p} = \begin{cases} \Delta', \tilde{s}^{\tilde{m}} \colon \{H_{\mathtt{q}}@\mathtt{q}\}_{\mathtt{q}\in I\setminus\{\mathtt{p}\}} \cup \{H_{\mathtt{p}}; M@\mathtt{p}\} \\ \qquad \text{if } \mathtt{p} \in I \ \ \Delta = \Delta', \tilde{s}^{\tilde{m}} \colon \{H_{\mathtt{q}}@\mathtt{q}\}_{\mathtt{q}\in I} \\ \Delta', \tilde{s}^{\tilde{m}} \colon \{H_{\mathtt{q}}@\mathtt{q}\}_{\mathtt{q}\in I} \cup \{M@\mathtt{p}\} \\ \qquad \text{if } \mathtt{p} \notin I \ \ \Delta = \Delta', \tilde{s}^{\tilde{m}} \colon \{H_{\mathtt{q}}@\mathtt{q}\}_{\mathtt{q}\in I} \\ \Delta', \tilde{s}^{\tilde{m}} \colon \{M@\mathtt{p}\} \\ \qquad \text{otherwise} \end{cases}$$

Fig. 1 Typing system for initial processes

$$\Gamma, a : S \vdash a : S \qquad \Gamma \vdash \mathsf{true}, \mathsf{false} : \mathsf{bool} \qquad \frac{\Gamma \vdash e_i \triangleright \mathsf{bool}}{\Gamma \vdash e_1 \mathtt{or}\ e_2 : \mathsf{bool}} \qquad \text{[NAME], [BOOL], [OR]}$$

$$\frac{\Gamma \vdash a : G \quad \Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : (G \upharpoonright 1)@1 \quad |\tilde{s}| = chans(G) \quad \mathscr{B}_k\langle G \rangle = m_k}{\Gamma \vdash \overline{a}_{[2..n]}(\tilde{s}^{\tilde{m}}).P \triangleright \Delta} \qquad \text{[MCAST]}$$

$$\frac{\Gamma \vdash a : G \quad \Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : (G \upharpoonright \mathtt{p})@\mathtt{p} \quad |\tilde{s}| = |\mathsf{sid}(G)|}{\Gamma \vdash a_{[p]}(\tilde{s}).P \triangleright \Delta} \qquad \text{[MACC]}$$

$$\frac{\Gamma \vdash e : S \qquad \Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : T@\mathtt{p}}{\Gamma \vdash s_k!\langle e \rangle; P \triangleright \Delta, \tilde{s}^{\tilde{m}} : k!\langle S \rangle; T@\mathtt{p}} \qquad \text{[SEND]}$$

$$\frac{\Gamma, x : S \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : T@\mathtt{p}}{\Gamma \vdash s_k?(x); P \triangleright \Delta, \tilde{s}^{\tilde{m}} : k?\langle S \rangle; T@\mathtt{p}} \qquad \text{[RCV]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : T@\mathtt{p}}{\Gamma \vdash s_k!\langle\!\langle \tilde{t} \rangle\!\rangle; P \triangleright \Delta, \tilde{s}^{\tilde{m}} : k!\langle T'@\mathtt{p}' \rangle; T@\mathtt{p},\ \tilde{t}^{\tilde{n}} : T'@\mathtt{p}'} \qquad \text{[DELEG]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : T@\mathtt{p},\ \tilde{t}^{\tilde{n}} : T'@\mathtt{p}'}{\Gamma \vdash s_k?(\!(\tilde{t})\!); P \triangleright \Delta, \tilde{s}^{\tilde{m}} : k?\langle T'@\mathtt{p}' \rangle; T@\mathtt{p}} \qquad \text{[SREC]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta, \tilde{s}^{\tilde{m}} : T_j@\mathtt{p} \qquad j \in I}{\Gamma \vdash s_k \lhd l_j P \triangleright \Delta, \tilde{s}^{\tilde{m}} : k \oplus \{l_i : T_i\}_{i \in I}@\mathtt{p}} \qquad \text{[SEL]}$$

$$\frac{\Gamma \vdash P_i \triangleright \Delta, \tilde{s}^{\tilde{m}} : T_i@\mathtt{p} \qquad \forall i \in I}{\Gamma \vdash\ s_k \rhd \{l_i : P_i\}_{i \in I} \triangleright \Delta, \tilde{s}^{\tilde{m}} : k\,\&\{l_i : T_i\}_{i \in I}@\mathtt{p}} \qquad \text{[BRA]}$$

$$\frac{\Gamma \vdash P \triangleright \Delta \qquad \Gamma \vdash Q \triangleright \Delta'}{\Gamma \vdash P \mid Q \triangleright \Delta, \Delta'} \qquad \text{[CONC]}$$

$$\frac{\Gamma \vdash e \triangleright \mathsf{bool} \qquad \Gamma \vdash P \triangleright \Delta \qquad \Gamma \vdash Q \triangleright \Delta}{\Gamma \vdash \mathsf{if}\ e\ \mathsf{then}\ P\ \mathsf{else}\ Q \triangleright \Delta} \qquad \text{[IF]}$$

$$\frac{\Delta\ \mathsf{end\ only}}{\Gamma \vdash \mathbf{0} \triangleright \Delta} \qquad \frac{\Gamma, a : G \vdash P \triangleright \Delta}{\Gamma \vdash (\nu a)P \triangleright \Delta} \qquad \text{[INACT],[NRES]}$$

$$\frac{\Delta\ \mathsf{end\ only}}{\Gamma, X : \Delta \vdash X \triangleright \Delta,} \qquad \frac{\Gamma, X : \Delta \vdash P \triangleright \Delta}{\Gamma \vdash \mu X.P \triangleright \Delta} \qquad \text{[VAR],[DEF]}$$

When we type parallel composition of processes, we need to calculate the resulting buffer sizes and check that they do not exceed the specified sizes. The parallel composition of session environments, denoted by $\Delta_1 \circ \Delta_2$, is defined as:

$$\begin{array}{l} \Delta \setminus \mathsf{dom}(\Delta') \cup \Delta' \setminus \mathsf{dom}(\Delta) \\ \cup \{\tilde{s}^{\tilde{n}} : \{H_\mathtt{p} \circ H'_\mathtt{p}@\mathtt{p}\}_{\mathtt{p} \in I} \mid \tilde{s}^{\tilde{n}} : \{H_\mathtt{p}@\mathtt{p}\}_{\mathtt{p} \in I} \in \Delta, \tilde{s}^{\tilde{n}} : \{H'_\mathtt{p}@\mathtt{p}\}_{\mathtt{p} \in I} \in \Delta', \mathscr{B}_k\langle\{H_\mathtt{p} \circ H'_\mathtt{p}@\mathtt{p}\}_{\mathtt{p} \in I}\rangle \leq n_k\} \end{array}$$

Fig. 2 Typing rules for buffers

$$\frac{}{\Gamma \vdash_{\{s_k\}} s_k : \emptyset \triangleright \emptyset} \qquad \frac{\Gamma \vdash v_i : S_i \qquad \Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \triangleright \Delta}{\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \cdot \tilde{v} \triangleright \Delta; \tilde{s}^{\tilde{m}} : k!\langle \tilde{S} \rangle @ \mathtt{p}} \quad [\textsc{Qnil}],[\textsc{Qval}]$$

$$\frac{\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \triangleright \Delta}{\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \cdot \tilde{t} \triangleright \Delta; \tilde{s}^{\tilde{m}} : k!\langle T @ \mathtt{q} \rangle @ \mathtt{p}, \tilde{t} : T @ \mathtt{q}} \quad [\textsc{Qsess}]$$

$$\frac{\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \triangleright \Delta}{\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \cdot l \triangleright \Delta; \tilde{s}^{\tilde{m}} : k \oplus l @ \mathtt{p}} \quad [\textsc{Qsel}]$$

where $H \circ H'$ is defined as $H;H'$ if H is a message type; or $H';H$ if H' is a message type; and otherwise undefined. The operator $\circ$ on environments Δ, Δ' is defined if the result respects the buffer bounds. The interesting two rules for runtime processes which contain buffers are:

$$\frac{\Gamma \vdash_{\Sigma} P \triangleright \Delta \qquad \Gamma \vdash_{\Sigma'} Q \triangleright \Delta' \qquad \Sigma \cap \Sigma' = \emptyset}{\Gamma \vdash_{\Sigma \cup \Sigma'} P \mid Q \triangleright \Delta \circ \Delta'}$$

$$\frac{\Gamma \vdash_{\Sigma} P \triangleright \Delta, \tilde{s}^{\tilde{n}} : \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \qquad \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \text{ coherent} \quad \mathscr{B}_k\langle \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \rangle \leq n_k \qquad \tilde{s} \in \Sigma}{\Gamma \vdash_{\Sigma \setminus \tilde{s}} (\nu\, \tilde{s}) P \triangleright \Delta}$$

The judgement $\Gamma \vdash_{\Sigma} P \triangleright \Delta$ means that P contains the buffers whose session names are in Δ. We define the full rules in figure 3.

Fig. 3 Typing rules for runtime processes

$$\frac{\Gamma \vdash P \triangleright \Delta}{\Gamma \vdash_{\emptyset} P \triangleright \Delta} \qquad \frac{\Gamma \vdash_{\Sigma} P \triangleright \Delta \qquad \Delta' \text{end only}}{\Gamma \vdash_{\Sigma} P \triangleright \Delta \circ \Delta'} \quad [\textsc{Dnil}], [\textsc{Dweak}]$$

$$\frac{\Gamma \vdash_{\Sigma} P \triangleright \Delta \qquad \Gamma \vdash_{\Sigma'} Q \triangleright \Delta' \qquad \Sigma \cap \Sigma' = \emptyset}{\Gamma \vdash_{\Sigma \cup \Sigma'} P \mid Q \triangleright \Delta \circ \Delta'} \quad [\textsc{Dconc}]$$

$$\frac{\Gamma \vdash_{\Sigma} P \triangleright \Delta, \tilde{s}^{\tilde{n}} : \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \quad \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \text{ coherent} \quad \mathscr{B}_k\langle \{T_{\mathtt{p}} @ \mathtt{p}\}_{\mathtt{p} \in I} \rangle \leq n_k \quad \tilde{s} \in \Sigma}{\Gamma \vdash_{\Sigma \setminus \tilde{s}} (\nu\, \tilde{s}) P \triangleright \Delta} \quad [\textsc{Dres}]$$

$$\frac{\Gamma, a : G \vdash_{\Sigma} P \triangleright \Delta}{\Gamma \vdash_{\Sigma} (\nu\, a) P \triangleright \Delta} \qquad \frac{\Gamma, X : \Delta \vdash P \triangleright \Delta}{\Gamma \vdash_{\Sigma} \mu X.P \triangleright \Delta} \quad [\textsc{Dname}],[\textsc{Dvar}]$$

For the parallel composition, we use $\circ$ and make sure that all buffers are disjoint. In the rule [DConc], we check that, when queues are composed, the family at $\tilde{s}$ is coherent (i.e. satisfies the linearity constraint [10]), and does not exceed the upper bounds of the buffers.

To prove subject reduction, we keep track of the correspondence between the session environment and the size of the buffer. We use the reduction over session typing,

$\Delta \xrightarrow{k} \Delta'$ generated by the following rules:

$$k!\langle U\rangle;H@\mathtt{p},k?\langle U\rangle;H'@\mathtt{q} \xrightarrow{k} H@\mathtt{p},H'@\mathtt{q}$$
$$k\oplus;\{l:H,...\}@\mathtt{p} \xrightarrow{k} k\oplus l;H@\mathtt{p}$$
$$k\oplus l;H@\mathtt{p},k\&\{l:H',...\}@\mathtt{q} \xrightarrow{k} H@\mathtt{p},H'@\mathtt{q}$$
$$\Delta \xrightarrow{s} \Delta' \implies \Delta,\Delta_0 \xrightarrow{s} \Delta',\Delta_0$$
$$H_1@\mathtt{p}_1,H_2@\mathtt{p}_2 \xrightarrow{k} H_1'@\mathtt{p}_1,H_2'@\mathtt{p}_2 \implies \tilde{s}\colon\{H_1@\mathtt{p}_1,H_2@\mathtt{p}_2,...\}_{i\in I} \xrightarrow{s_k} \tilde{s}\colon\{H_1'@\mathtt{p}_1,H_2'@\mathtt{p}_2,...\}_{i\in I}$$

C.4 Appendix for Proofs of Subject Reduction Theorem and Buffer Safety

This subsection gives the proofs of the subject reduction theorem and buffer safety.

Below we often write $\mathscr{B}_k\langle G\rangle$ and $[\![G]\!] = \{T_\mathtt{p}@\mathtt{p}\}_{\mathtt{p}\in I}$ as $\mathscr{B}_k\langle\{T_\mathtt{p}@\mathtt{p}\}_{\mathtt{p}\in I}\rangle$.

Proof ***(Proofs of Lemma 3.7).*** Mechanical by case analysis on the environment reduction rules.

Lemma C.2 *$\Gamma \vdash_\Sigma P \triangleright \Delta$ and $P \equiv Q$ imply $\Gamma \vdash_\Sigma P \triangleright \Delta$.*

Proof. By induction on the derivation $P \equiv Q$. We assume in each case that $\Gamma \vdash_\Sigma P \triangleright \Delta$.

Case $P \mid \mathbf{0} \equiv P$. Assuming that $\Gamma \vdash_\Sigma P \triangleright \Delta$ and knowing that $\Gamma \vdash_\emptyset \mathbf{0} \triangleright \emptyset$ by [DNIL], we can use rule [DCONC] to get the desired $\Gamma \vdash_\Sigma P \mid \mathbf{0} \triangleright \Delta$. In the other direction, we know that $\Gamma \vdash_\Sigma P \mid \mathbf{0} \triangleright \Delta$. The last rule is [DCONC]. We thus have $\Delta = \Delta_1 \circ \Delta_2$ and $\Sigma = \Sigma_1 \uplus \Sigma_2$ with $\Gamma \vdash_{\Sigma_1} P \triangleright \Delta_1$ and $\Gamma \vdash_{\Sigma_2} \mathbf{0} \triangleright \Delta_2$. By [INACT], we have $\Sigma_2 = \emptyset$ and Δ_2 has only end. We can then use weakening to get the desired $\Gamma \vdash_{\Sigma_1} P \triangleright \Delta_1 \circ \Delta_2$.

Case $P \mid Q \equiv Q \mid P$. Assuming $\Gamma \vdash_\Sigma P \mid Q \triangleright \Delta$, we get from [DCONC] the existence of $\Delta = \Delta_1 \circ \Delta_2$ and $\Sigma = \Sigma_1 \uplus \Sigma_2$ such that $\Gamma \vdash_{\Sigma_1} P \triangleright \Delta_1$ and $\Gamma \vdash_{\Sigma_2} Q \triangleright \Delta_2$. Since the $\circ$ and $\uplus$ are commutative, we can apply [DCONC] with permuted premises to get $\Gamma \vdash_\Sigma Q \mid P \triangleright \Delta$.

Case $(P \mid Q) \mid R \equiv P \mid (Q \mid R)$. Similar to the previous case, but using the associativity of $\uplus$ and $\circ$ instead of the commutativity.

Case $(\nu n)P \mid Q \equiv (\nu n)(P \mid Q)$ if $n \notin \mathsf{fn}(Q)$. Assuming that $\Gamma \vdash_\Sigma (\nu n)P \mid Q \triangleright \Delta$, we get from [DCONC] the existence of $\Delta = \Delta_1 \circ \Delta_2$ and $\Sigma = \Sigma_1 \uplus \Sigma_2$ such that $\Gamma \vdash_{\Sigma_1} (\nu n)P \triangleright \Delta_1$ and $\Gamma \vdash_{\Sigma_2} Q \triangleright \Delta_2$. We invert [DNAME] and conclude by weakening on the typing judgement of Q, and apply [DCONC] then [DNAME]. The other direction relies on the strenghtening of the typing judgement with a name absent from the free names of Q. **Case $(\nu nn')P \equiv (\nu n'n)P$.** By the reordering of environment bindings.

Case $(\nu n)\mathbf{0} \equiv \mathbf{0}$. Trivial.

Case $\mu X.P \equiv P[\mu X.P/X]$. Assuming that $\Gamma \vdash_\Sigma \mu X.P \triangleright \Delta$, we get from [DVAR] that $\Gamma, X : \Delta \vdash_\Sigma P \triangleright \Delta$. We conclude by type preservation by substitution of process variable.

Case $(\nu s_1..s_n)\Pi_i s_i^{n_i} : \emptyset \equiv \mathbf{0}$. We repeatedly use [QNIL], [DCONC] and [DNAME].

Proof ***(Proof of Theorem 3.8).***
$\Gamma \vdash_\Sigma P \triangleright \Delta$ and $P \longrightarrow Q$ with fully coherent $(\Gamma, \Sigma, P, \Delta)$ imply $\Gamma \vdash_\Sigma Q \triangleright \Delta'$ for some Δ', s_k such that $\Delta = \Delta'$ or $\Delta(\xrightarrow{s_k})^*\Delta'$ and $\mathscr{B}_k\langle\Delta(\tilde{s})\rangle \geq \mathscr{B}_k\langle\Delta'(\tilde{s})\rangle$, with $(\Gamma, \Sigma, Q, \Delta')$ fully coherent. The property $\mathscr{B}_k\langle\Delta(\tilde{s})\rangle \geq \mathscr{B}_k\langle\Delta'(\tilde{s})\rangle$ comes from lemma 3.7.

By induction on the derivation $P \longrightarrow Q$.

Case [LINK]**.** Let $P = \overline{a}_{[2..\mathtt{n}]}(\tilde{s}^{\tilde{m}}).P_1 \mid a_{[2]}(\tilde{s}).P_2 \mid \cdots \mid a_{[\mathtt{n}]}(\tilde{s}).P_\mathtt{n}$ and $Q = (\nu\, \tilde{s})(P_1 \mid P_2 \mid \ldots \mid P_\mathtt{n} \mid s_1^{n_1} : \emptyset \mid \cdots \mid s_m^{n_m} : \emptyset)$. We assume that $\Gamma \vdash_\Sigma P \triangleright \Delta$. We know that $a \in \Gamma$. The last rule is [CONC]. We thus have $\Delta = \Delta_1 \circ \ldots \circ \Delta_\mathtt{n}$ and $\Sigma = \emptyset$ such that $\Gamma \vdash_\emptyset \overline{a}_{[2..\mathtt{n}]}(\tilde{s}^{\tilde{m}}).P_1 \triangleright \Delta_1$ and $\Gamma \vdash_\emptyset a_{[\mathtt{i}]}(\tilde{s}).P_\mathtt{i} \triangleright \Delta_\mathtt{i}$ for $2 \leq i \leq n$. By reversing [MCAST] and [MACC], we get $\Gamma \vdash_\emptyset P_\mathtt{i} \triangleright \Delta_\mathtt{i}, \tilde{s}^{\tilde{m}} : (G \upharpoonright \mathtt{i})@\mathtt{i}$ for $1 \leq i \leq n$. From repeated applications of [QNIL] and [DCONC], we get $\Gamma \vdash_{\tilde{s}} \Pi_i P_\mathtt{i} \mid \Pi_i s_i : \emptyset \triangleright \Delta, \{\tilde{s}^{\tilde{m}} : (G \upharpoonright \mathtt{i})@\mathtt{i}\}_{i \leq n}$. We conclude by [DRES].

Case [SEND]**.** Let $P = s!\langle e\rangle; R \mid s^n : \tilde{h} \cdot v$ and $Q = R \mid s^n : \tilde{h} \cdot v$ with $n \gtrsim |\tilde{h}|$, $e \downarrow v$. We assume that $\Gamma \vdash_\Sigma P \triangleright \Delta$. By reversing typing rules [DCONC], we know that $\Gamma \vdash_\emptyset s!\langle e\rangle; R \triangleright \Delta_1$ and that $\Gamma \vdash_{s^n} s^n : \tilde{h} \triangleright \Delta_2$ with $\Delta = \Delta_1 \circ \Delta_2$ and $\Sigma = \{s^n\}$. By reversing [SEND], we have $\Gamma \vdash e : S$ and $\Delta_1 = \Delta_1', \tilde{s}^{\tilde{m}} : k!\langle S\rangle; T@\mathtt{p}$ with $\Gamma \vdash_\emptyset R \triangleright \Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}$. By [QVAL] and type preservation for expression reduction, we get $\Gamma \vdash_{\{s_k\}} s_k : \tilde{h} \cdot \tilde{v} \triangleright \Delta_2; \tilde{s}^{\tilde{m}} : k!\langle\tilde{S}\rangle@\mathtt{p}$. By [DCONC] we get $\Gamma \vdash_\Sigma Q \triangleright (\Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}) \circ (\Delta_2; \tilde{s}^{\tilde{m}} : k!\langle\tilde{S}\rangle@\mathtt{p})$. Since the definition of $\circ$ gives us $(\Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}) \circ (\Delta_2; \tilde{s}^{\tilde{m}} : k!\langle\tilde{S}\rangle@\mathtt{p}) = (\Delta_1', \tilde{s}^{\tilde{m}} : k!\langle S\rangle; T@\mathtt{p}) \circ (\Delta_2) = \Delta$, we can conclude that $\Gamma \vdash_\Sigma Q \triangleright \Delta$.

Case [SEL]**.** Similar to [SEND].

Case [RECV]**.** Let $P = s?(x); R \mid s^n : v \cdot \tilde{h}$ and $Q = R[v/x] \mid s^n : \tilde{h}$. We assume that $\Gamma \vdash_\Sigma P \triangleright \Delta$. By reversing typing rules [DCONC], we know that $\Gamma \vdash_\emptyset s?(x); R \triangleright \Delta_1$ and that $\Gamma \vdash_{s^n} s^n : v \cdot \tilde{h} \triangleright \Delta_2$ with $\Delta = \Delta_1 \circ \Delta_2$ and $\Sigma = \{s^n\}$. By reversing [RCV], $\Delta_1 = \Delta_1', \tilde{s}^{\tilde{m}} : k?\langle S\rangle; T@\mathtt{p}$ with $\Gamma, x : S \vdash_\emptyset R \triangleright \Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}$. By type preservation for expression substitution, we have $\Gamma \vdash_\emptyset R[v/x] \triangleright \Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}$. By queue typing reversion, we have $\Delta_2 = \Delta_2', \tilde{s}^{\tilde{m}} : k!\langle S\rangle; T'@\mathtt{q}$ with $\Gamma \vdash_{s^n} s^n : \tilde{h} \triangleright \Delta_2', \tilde{s}^{\tilde{m}} : T'@\mathtt{q}$ and $\Gamma \vdash v : S$. By the definition of $\circ$, we know that $p \neq q$. Then we have the reduction $\Delta = \Delta_1 \circ \Delta_2 = (\Delta_1', \tilde{s}^{\tilde{m}} : k?\langle S\rangle; T@\mathtt{p}) \circ (\Delta_2', \tilde{s}^{\tilde{m}} : k!\langle S\rangle; T'@\mathtt{q}) \xrightarrow{k} (\Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}) \circ (\Delta_2', \tilde{s}^{\tilde{m}} : T'@\mathtt{q})$. By [DCONC], we get $\Gamma \vdash_{s^n} R[v/x] \mid s^n : \tilde{h} \triangleright (\Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}) \circ (\Delta_2', \tilde{s}^{\tilde{m}} : T'@\mathtt{q})$. In the "coherent for s_k" environment, we can conclude from lemma 3.7.

Case [BRA]**.** Similar to [RECV].

Case [DELEG]**.** Let $P = s!\langle\!\langle\tilde{t}\rangle\!\rangle; R \mid s : \tilde{h}$ and $Q = R \mid s^n : \tilde{h} \cdot \tilde{t}$ with $n \gtrsim |\tilde{h}|$. We assume that $\Gamma \vdash_\Sigma P \triangleright \Delta$. We proceed in a similar way to [SEND]: we just use [QSESS] instead of [QVAL].

Case [SREC]**.** Let $P = s?(\!(\tilde{t})\!); R \mid s^n : \tilde{t} \cdot \tilde{h}$ and $Q = R \mid s^n : \tilde{h}$. We assume that $\Gamma \vdash_\Sigma P \triangleright \Delta$. By reversing typing rules [DCONC], we know that $\Gamma \vdash_\emptyset s?(\!(\tilde{t})\!); R \triangleright \Delta_1$ and that $\Gamma \vdash_{s^n} s^n : v \cdot \tilde{h} \triangleright \Delta_2$ with $\Delta = \Delta_1 \circ \Delta_2$ and $\Sigma = \{s^n\}$. By reversing [SREC], $\Delta_1 = \Delta_1', \tilde{s}^{\tilde{m}} : k?\langle T'@\mathtt{p}'\rangle; T@\mathtt{p}$ with $\Gamma \vdash R \triangleright \Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}, \tilde{t}^{\tilde{n}} : T'@\mathtt{p}'$. By queue typing reversion, we have $\Delta_2 = \Delta_2', \tilde{s}^{\tilde{m}} : k!\langle T@\mathtt{q}\rangle; T'@\mathtt{p}, \tilde{t} : T@\mathtt{q}; T''$. By the definition of $\circ$, we know that $p \neq q$. Then we have the reduction $\Delta = \Delta_1 \circ \Delta_2 = (\Delta_1', \tilde{s}^{\tilde{m}} : k?\langle T'@\mathtt{p}'\rangle; T@\mathtt{p}) \circ (\Delta_2', \tilde{s}^{\tilde{m}} : k!\langle T@\mathtt{q}\rangle; T'@\mathtt{p}, \tilde{t} : T@\mathtt{q}; T'') \xrightarrow{k} (\Delta_1', \tilde{s}^{\tilde{m}} : T@\mathtt{p}) \circ (\Delta_2', \tilde{s}^{\tilde{m}} : T'@\mathtt{p}, \tilde{t} : T@\mathtt{q}; T'')$. We conclude by [DCONC].

For the remaining cases, it is sufficient to give a sketch proof since they are standard.

Case [IFT] and [IFF]**.** Trivial.

Case [SCOP]**.** By induction with [NRES].

Case [PAR]**.** By induction with [DCONC].

Case [DEFIN]**.** By induction with [DEF].

Case [STR]**.** By lemma C.2.

*Proof (**Proofs of Buffer Safety (Corollary 3.9)).*** By Subject Congruence and Reduction, it is sufficient to prove: If $\Gamma \vdash_\Sigma P \triangleright \Delta$ then $P \not\to \mathsf{Err}$. Suppose $P = s!\langle e\rangle; Q \mid s^n : \tilde{h}$ with $n \leq |\tilde{h}|$. This contradicts the typing system of the parallel composition [DCONC] since $\circ$ is not defined. The remaining cases (selection and delegation) are similar. Other cases are by inductive hypothesis.

D Appendix for Section 4: Channel Attribution

We give here the omitted definitions and proofs, along with additional details from Section 4.

Definition D.1 (stripped session type) We formally define the *stripped session types* $\underline{G}$ by:

$$\begin{array}{rcl} \underline{G} ::= & \mathtt{p} \to \mathtt{p}'\langle U\rangle;\underline{G}' \text{ values} & \mid \mathtt{p} \to \mathtt{p}'\{l_j\colon \underline{G}_j\}_{j\in J} \text{ branching} \\ \mid & \mu \mathbf{x}.\underline{G} \mid \mathbf{x} \quad \text{recursion} & \mid \mathsf{end} \qquad\qquad \text{end} \end{array}$$

*Proof (**Proof of Lemma 4.1).***

1. Since in the singleton allocation, two prefixes using the same channel always involve the same sender and receiver, there is always an I-dependency and an O-dependencies between these prefixes.
2. We rely on the fact that if the channel is finite, there is an IO-dependency between its only two uses in any cycle of the unfolded graph. Therefore, the reset property holds when the bound computation algorithm checks it: the bound is exactly 1.
3. The total number of channels is bounded by the size of the global type.

Lemma D.2 *The existence of an infinite channel in a session does not depend on any particular set of channel equalities. If $\mathscr{B}_k\langle G\rangle = \infty$ then $\forall E, \mathscr{B}_k\langle G_E\rangle = \infty$.*

*Proof (**Proof of Lemma D.2).*** Since each equality embeds a linearity check, and that linearity imposes a II and OO dependency between uses of the same channel, it follows that an IO-dependency cannot be broken by a channel equality.

Branch allocation. The branch allocation is a simple refinement of the singleton allocation.

Definition D.3 (branch allocation) *Singleton allocations can be optimised by allowing the sharing of channels between branches that are not under a recursion operator.*

The branch allocation preserves linearity and is efficient to compute. The sizes of the channels in the branch allocation verify the same properties as in the singleton allocation.

Principal allocation. The most widely used allocation method attributes two communication channels (one in each direction) for each pair of participants.

Proof ***(Proof of Lemma 4.3).***

1. By the same argument as in lemma 4.1(1).
2. There is at most one channel per pair of participants.

E Appendix for Section 5.1: Confirmation Message Insertion

The unfolding function unfolds under the recursion and not outside. The unfolding function is defined by:

$$\begin{aligned}
\phi^1(\mathtt{p} \to \mathtt{p}' \colon k\,\langle U\rangle; G') &= \mathtt{p} \to \mathtt{p}' \colon k\,\langle U\rangle; \phi^1(G') \\
\phi^1(\mathtt{p} \to \mathtt{p}' \colon k\,\{l_j \colon G_j\}_{j\in J}) &= \mathtt{p} \to \mathtt{p}' \colon k\,\{l_j \colon \phi^1(G_j)\}_{j\in J} \\
\phi^1(\mu \mathbf{x}.G) &= \mu \mathbf{x}.\phi^1(G[\mu \mathbf{x}^*.G/\mathbf{x}]) \\
\phi^1(\mu \mathbf{x}^*.G) &= G \\
\phi^1(\mathsf{end}) &= \mathsf{end}
\end{aligned}$$

The additional proposition about the algorithms given below.

F Appendix for Section 5.2: Messaging Optimisations

This section gives the omitted definitions from Section 5.2. The global types and local types are extended with annotations by a set of nodes N.

$$\begin{aligned}
G &::= \mathtt{p} \to \mathtt{p}' \colon k^N\,\langle U\rangle; G' \mid \mathtt{p} \to \mathtt{p}' \colon k^N\,\{l_j \colon G_j\}_{j\in J} \mid \cdots \\
T &::= k^N!\langle U\rangle; T \mid k^N?\langle U\rangle; T \mid k^N \oplus \{l_i \colon T_i\}_{i\in I} \mid k^N \& \{l_i : T_i\}_{i\in I} \mid \cdots
\end{aligned}$$

We often omit N if it is not necessary. The projection is defined just as before on the extended types.

Definition F.1 (action subtyping) *We assume $k \neq k_0$.*

$$\begin{array}{lll}
(\mathsf{OI}) & k^N!\langle U\rangle; k_0^{N_0}?\langle U'\rangle; T \ll k_0^{N_0}?\langle U'\rangle; k^N!\langle U\rangle; T & N_0 \cap N = \emptyset \\
(\mathsf{OB}) & k^N!\langle U\rangle; k_0^{N_0} \& \{l_j : T_j\}_{j\in J} \ll k_0^{N_0} \& \{l_j : k^N!\langle U\rangle; T_j\}_{j\in J} & N_0 \cap N = \emptyset \\
(\mathsf{SI}) & k^N \oplus \{l_j : k_0^{N_0}?\langle U\rangle; T_j\}_{j\in J} \ll k_0^{N_0}?\langle U\rangle; k^N \oplus \{l_j : T_j\}_{j\in J} & N_0 \cap N = \emptyset \\
(\mathsf{SB}) & k^N \oplus \{l_i : k_0^{N_0} \& \{l'_j : T_{ij}\}_{j\in J}\}_{i\in I} \ll k_0^{N_0} \& \{l'_j : k^N \oplus \{l_i : T_{ij}\}_{i\in I}\}_{j\in J} & N_0 \cap N = \emptyset \\
(\mathsf{OO}) & k!\langle U\rangle; k_0!\langle U'\rangle; T \ll k_0!\langle U'\rangle; k!\langle U\rangle; T & \\
(\mathsf{II}) & k?\langle U\rangle; k_0?\langle U'\rangle; T \ll k_0?\langle U'\rangle; k?\langle U\rangle; T & \\
(\mathsf{SO}) & k \oplus \{l_i : k_0!\langle U\rangle; T_i\}_{i\in I} \ll k_0!\langle U\rangle; k \oplus \{l_i : T_i\}_{i\in I} & \\
(\mathsf{OS}) & k_0!\langle U\rangle; k \oplus \{l_i : T_i\}_{i\in I} \ll k \oplus \{l_i : k_0!\langle U\rangle; T_i\}_{i\in I} & \\
(\mathsf{SS}) & k \oplus \{l_i : k_0 \oplus \{l'_j : T_{ij}\}_{j\in J}\}_{i\in I} \ll k_0 \oplus \{l'_j : k \oplus \{l_i : T_{ij}\}_{i\in I}\}_{j\in J} &
\end{array}$$

$$(\mathsf{CO})\ \frac{T \ll T'}{k!\langle U\rangle;T \ll k!\langle U\rangle;T'} \qquad (\mathsf{CI})\ \frac{T \ll T'}{k?\langle U\rangle;T \ll k?\langle U\rangle;T'}$$

$$(\mathsf{CB})\ \frac{\forall i \in I.\ T_i \ll T_i'}{k\&\{l_i:T_i\}_{i\in I} \ll k\&\{l_i:T_i'\}_{i\in I}}$$

$$(\mathsf{CS})\ \frac{\forall i \in I.\ T_i \ll T_i'}{k\oplus\{l_i:T_i\}_{i\in I} \ll k\oplus\{l_i:T_i'\}_{i\in I}}$$

$$(\mathsf{Tr})\ \frac{T_1 \ll T_2 \quad T_2 \ll T_3}{T_1 \ll T_3} \qquad (\mathsf{E})\ \mathsf{end} \ll \mathsf{end} \qquad (\mathsf{M})\ \mu\mathbf{x}.T \ll \mu\mathbf{x}.T$$

We omit $(\mathsf{BI},\mathsf{IB},\mathsf{BB})$.

Definition F.2 (n-time unfolding) For clarity, we omit N from the definition below.

$$\begin{array}{ll}
\varphi^0(T) = T \text{ for all } T & \varphi^{1+n}(T) = \varphi^1(\varphi^n(T)) \\
\varphi^1(k!\langle U\rangle;T) = k!\langle U\rangle;\varphi^1(T) & \varphi^1(k\oplus\{l_i\colon T_i\}_{i\in I}) = k\oplus\{l_i\colon \varphi^1(T_i)\}_{i\in I} \\
\varphi^1(k?\langle U\rangle;T) = k?\langle U\rangle;\varphi^1(T) & \varphi^1(k\&\{l_i\colon T_i\}_{i\in I}) = k\&\{l_i\colon \varphi^1(T_i)\}_{i\in I} \\
\varphi^1(\mu\mathbf{x}.T) = T[\mu\mathbf{x}.T/\mathbf{x}] & \varphi^1(\mathbf{x}) = \mathbf{x} \\
\varphi^1(\mathsf{end}) = \mathsf{end} &
\end{array}$$

Below we define the asynchronous communication subtyping from [14]. We omit delegation for simplicity (see [13] for the full definition). Below we write *Type* for the collection of all closed local types.

Definition F.3 (Size-preserving asynchronous subtyping) A relation $\mathfrak{R} \in \mathit{Type} \times \mathit{Type}$ is an asynchronous type simulation if $(T_1,T_2) \in \mathfrak{R}$ implies the following conditions. This relation is not related to the type annotations.

- If $T_1 = \mathsf{end}$, then $\varphi^n(T_2) = \mathsf{end}$.
- If $T_1 = k!\langle U_1\rangle;T_1'$, then $\varphi^n(T_2) \gg k!\langle U_2\rangle;T_2'$, $(T_1',T_2') \in \mathfrak{R}$ and $(U_1,U_2) \in \mathfrak{R}$.
- If $T_1 = k?\langle U_1\rangle;T_1'$, then $\varphi^n(T_2) \gg k?\langle U_2\rangle;T_2'$, $(T_1',T_2') \in \mathfrak{R}$ and $(U_2,U_1) \in \mathfrak{R}$.
- If $T_1 = k\oplus\{l_i : T_{1i}\}_{i\in I}$, then $\varphi^n(T_2) \gg k\oplus\{l_j : T_{2j}\}_{j\in J}$ and $I \subseteq J$ and $\forall i \in I.(T_{1i},T_{2i}) \in \mathfrak{R}$.
- If $T_1 = k\&\{l_i : T_{1i}\}_{i\in I}$, then $\varphi^n(T_2) \gg k\&\{l_j : T_{2j}\}_{j\in J}$ and $J \subseteq I$ and $\forall j \in J.(T_{1j},T_{2j}) \in \mathfrak{R}$.
- If $T_1 = \mu\mathbf{x}.T$, then $(\varphi^1(T_1),T_2) \in \mathfrak{R}$.

where a type simulation of $(U_1,U_2) \in \mathfrak{R}$ is defined as a standard bisimulation (since U is invariant). The coinductive subtyping relation $T_1 \leqslant_c T_2$ (read: T_1 is an *size-preserving asynchronous subtype* of T_2) is defined when there exists a type simulation $\mathfrak{R}$ with $(T_1,T_2) \in \mathfrak{R}$.

Lemma F.4 *Suppose* $[\![G]\!] = \{T@\mathtt{p}\}_{\mathtt{p}}$ *with* $T@\mathtt{p} = (G\upharpoonright\mathtt{p})@\mathtt{p}$ *and* $\mathtt{p} \in G$. *Assume* $T@\mathtt{p} \leqslant_c T'@\mathtt{p}$ *with* $G' = \{T'@\mathtt{p}\}_{\mathtt{p}}$ *Then* $\mathscr{B}_k\langle G\rangle = \mathscr{B}_k\langle G'\rangle$.

Proof. Obvious since the permutations do not alter the IO-ordering in $[\![G]\!]$ by definition.

Fig. 4 Double-Buffering

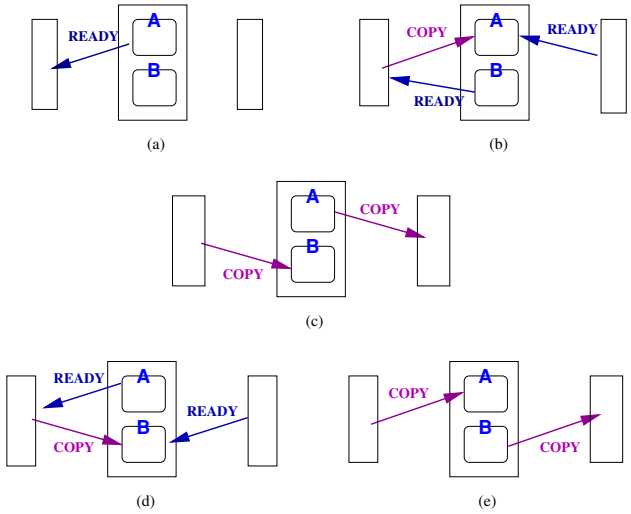

G Appendix for Section 6: Application Examples

G.1 Appendix for the Multi-Buffering Examples

We give additional explanations concerning the buffering algorithm, taking the case $n = 2$. It is sufficient to understand the basic mechanism from the double-buffering version of the algorithm. We take a simple stream program for data encryption explained by the following figure.

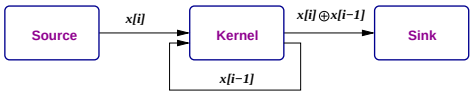

A data producer *Source* continuously feeds data to *Kernel*, which calculates the XOR of each element with a key and outputs the result to a stream to a consumer *Sink*. Kernel uses two arrays, or *buffers*, named A and B in the picture: while Source uses a single 16k array (in practice it can use a large cyclic buffer), fed by, say, a byte stream from an external channel. While Kernel is receiving data into the array A from Source, it processes data in the array B and sends the result to Sink, then repeats the same exchanging the roles of A and B. Figure 4 explains the following five steps:

(a) Kernel tells Source it is ready to receive an initial strip at A;
(b) Source immediately does so: asynchronously Kernel continues to tell Source it's *also* ready at B, and again asynchronously Sink tells Kernel it's ready to receive at its own array;

(c) Kernel finishes its processing at its A-strip and sends the resulting data to Sink, while Source is sending its strip to B at Kernel;
(d) The sending at B by Source continues, but (since Kernel has now sent out its A-strip) Kernel asynchronously tells Source it is ready at A; again asynchronously Sink tells the same to Kernel;
(e) Now the situation is symmetric to **(c)**: Source writing *to* A to Kernel and Kernel writing *from* B to Sink. We now go back to **(b)**.

The algorithm allows asynchrony among local computations and communications with minimal synchronisation to prevent data pollution, overlapping computation and communication, cf. [15, Double-buffering] and [5].

Proof. ***Proof of Proposition 6.1*** First we define the global type. For simplicity, we consider the double-buffering.

$$\begin{array}{l}\mu\mathbf{x}.(\ \mathsf{K}\to\mathsf{So}:r_0^{\mathsf{n}_{r_0},\mathsf{n}_{s_0}}\langle\rangle;\ \mathsf{So}\to\mathsf{K}:s_0^{\mathsf{n}_{r_0},\mathsf{n}_{s_0}}\langle U\rangle;\\ \quad\mathsf{Si}\to\mathsf{K}:\ t_0^{\mathsf{n}_{u_0},\mathsf{n}_{t_0}}\langle\rangle;\ \mathsf{K}\to\mathsf{Si}:\ u_0^{\mathsf{n}_{u_0},\mathsf{n}_{t_0}}\langle U\rangle;\\ \quad\mathsf{K}\to\mathsf{So}:r_1^{\mathsf{n}_{r_1},\mathsf{n}_{s_1}}\langle\rangle;\ \mathsf{So}\to\mathsf{K}:s_1^{\mathsf{n}_{r_1},\mathsf{n}_{s_1}}\langle U\rangle;\\ \quad\mathsf{Si}\to\mathsf{K}:\ t_1^{\mathsf{n}_{u_1},\mathsf{n}_{t_1}}\langle\rangle;\ \mathsf{K}\to\mathsf{Si}:\ u_1^{\mathsf{n}_{u_1},\mathsf{n}_{t_1}}\langle U\rangle;\mathbf{x})\end{array}$$

Thus we can permute the first outputs r_0 and r_1 at Kernel since they are disjoint with the previous prefixes. However the second output at r_0 (the second output created by the second unfolding) cannot be permuted with the first inputs at s_0 since r_0's set is not disjoint with s_0's set. Similarly for r_1 and s_1.

The communication safety and progress are respectively derived from Theorem 5.5 and Theorem 5.12 in [10].

By the above annotation, the following Unsafe Kernel (which leads to two units bounds) is not typable since it permutes the second r_1 with s_1.

$$\begin{array}{l}\textit{Unsafe Kernel: } r_0!\langle\rangle;r_1!\langle\rangle;\mu X.(s_0?(x_A);t_0?();u_0!\langle x_A\rangle;\\ \qquad\qquad r_0!\langle\rangle;r_1!\langle\rangle;s_1?(x_B);t_1?();u_1!\langle x_B\rangle;X)\end{array}$$

G.2 Appendix for MPSoC Buffer Allocations

This appendix gives the omitted definitions and some additional material about the MPSoC example from § 6

Remark G.1 (`foreach` **and sequencing)** As the readers might have already noticed, even when the two prefixes are syntactically composed sequentially in the global type, there might be no actual ordering. For example, in the following global type, there is no sequentiality imposed between $\{\text{Alice},\text{Bob}\}$ and $\{\text{Carol},\text{Dave}\}$.

$$\text{Alice}\to\text{Bob}:k\langle U\rangle;\text{Carol}\to\text{Dave}:k'\langle U\rangle$$

Hence the local processes typable from the above global type are also typable by the following global type:

$$\text{Carol}\to\text{Dave}:k'\langle U\rangle;\text{Alice}\to\text{Bob}:k\langle U\rangle$$

Similarly inside `foreach`, it is not needed to have explicit ordering. As such, we can think about global types up to asynchronous permutations (see [14]). In the MPSoC example, many actions can be permutable as explained below.

The local processes which confirm the projection of the following type is given below.

Original Global Type

$$\begin{array}{l}\mu \mathbf{x}.\ (\texttt{foreach}(i \in \{2,3\})\{\\ \quad \mathtt{p}[1] \to \mathtt{p}[i] : \left\{\begin{array}{l}\mathsf{on} : \mathtt{p}[i] \to \mathtt{p}[4] : \mathsf{on}; \mathtt{p}[4] \to \mathtt{p}[6] : \mathsf{on} \\ \mathsf{off} : \mathtt{p}[i] \to \mathtt{p}[5] : \mathsf{off}; \mathtt{p}[5] \to \mathtt{p}[6] : \mathsf{off}\end{array}\right\}\};\\ \quad \mathtt{p}[1] \to \mathtt{p}[6] : \langle \mathsf{real} \rangle); \mathbf{x}\end{array}$$

Local Processes We assume s_{ij} is a channel from $\mathtt{P}[i]$ to $\mathtt{P}[j]$.

$$\begin{array}{l}\mathtt{P}[1] = \mu X.s_{12} \lhd l; s_{13} \lhd l; s_{16}!\langle V \rangle; X\\ \mathtt{P}[2] = \mu X.s_{12} \rhd \{\mathsf{on} : s_{13} \lhd \mathsf{on}; X,\ \mathsf{off} : s_{25} \lhd \mathsf{off}; X\}\\ \mathtt{P}[3] = \mu X.s_{13} \rhd \{\mathsf{on} : s_{34} \lhd \mathsf{on}; X,\ \mathsf{off} : s_{35} \lhd \mathsf{off}; X\}\\ \mathtt{P}[4] = \mu X.s_{24} \rhd l_x; s_{34} \rhd l_y; s_{46}!\langle x+y \rangle; X\\ \mathtt{P}[5] = \mu X.s_{25} \rhd l_x; s_{35} \rhd l_y; s_{56}!\langle x+y \rangle; X\\ \mathtt{P}[6] = \mu X.s_{46}?(x); s_{56}?(y); s_{16}?(z); X\end{array}$$

where we use the flexible projection for branching from [21, § 3] and omits the useless branchings from $\mathtt{P}[4]$ and $\mathtt{P}[5]$.

Global Refinement: Optimal The first global refinement for the global type is given by the following type.

$$\mu \mathbf{x}.\texttt{foreach}(n < m)\{\texttt{foreach}(i \in \{2,3\})\{G[i]\}\}; \mathtt{p}[6] \to \mathtt{p}[1]\langle \mathsf{unit} \rangle : \mathbf{x}$$

As seen from the above type, the final confirmation imposes a global synchronisation from all participants after m times parallel iterations.

Local Processes The only different processes are $\mathtt{P}[1]$ and $\mathtt{P}[6]$. Below c_{61} is a channel for the confirmation message from $\mathtt{P}[6]$ to $\mathtt{P}[1]$.

$$\begin{array}{l}\mathtt{P}[1] = \mu X.\texttt{foreach}(n < m)\{s_{12} \lhd l_1; s_{13} \lhd l_3; s_{16} \lhd l_6\}; c_{61}?(); X\\ \mathtt{P}[6] = \mu X.\texttt{foreach}(n < m)\{s_{46}?(x); s_{56}?(y); s_{16}?(z)\}; c_{61}!\langle\rangle; X\end{array}$$

Global Refinement: Instant. The global type which is derived directly by applying the instant algorithm is given as below:

$$\begin{array}{l}\mu \mathbf{x}.\ \texttt{foreach}(n < m)\{\\ \quad \texttt{foreach}(i \in \{2,3\})\{\\ \quad \mathtt{p}[1] \to \mathtt{p}[i] : \left\{\begin{array}{l}\mathsf{on} : \mathtt{p}[i] \to \mathtt{p}[4] : \mathsf{on}, \mathtt{p}[4] \to \mathtt{p}[6] : \mathsf{on}, \\ \mathsf{off} : \mathtt{p}[i] \to \mathtt{p}[5] : \mathsf{off}; \mathtt{p}[5] \to \mathtt{p}[6] : \mathsf{off}, \end{array}\right\}\};\\ \quad \mathtt{p}[1] \to \mathtt{p}[6] : \langle \mathsf{real} \rangle;\\ \};\\ \mathtt{p}[2] \to \mathtt{p}[1] : \langle \mathsf{unit} \rangle; \mathtt{p}[3] \to \mathtt{p}[1] : \langle \mathsf{unit} \rangle; \mathtt{p}[4] \to \mathtt{p}[2] : \langle \mathsf{unit} \rangle;\\ \mathtt{p}[4] \to \mathtt{p}[3] : \langle \mathsf{unit} \rangle; \mathtt{p}[5] \to \mathtt{p}[2] : \langle \mathsf{unit} \rangle; \mathtt{p}[5] \to \mathtt{p}[3] : \langle \mathsf{unit} \rangle;\\ \mathtt{p}[6] \to \mathtt{p}[4] : \langle \mathsf{unit} \rangle; \mathtt{p}[6] \to \mathtt{p}[5] : \langle \mathsf{unit} \rangle; \mathtt{p}[6] \to \mathtt{p}[1] : \langle \mathsf{unit} \rangle; \mathbf{x}\end{array}$$

The causalities are the same as the following types, which gives the same end-point processes up to asynchronous permutations.

$$
\begin{array}{l}
\mu \mathbf{x}.\ (\texttt{foreach}(i \in \{2,3\})\{ \\
\quad \texttt{foreach}(n < m)\{ \\
\quad\quad \mathtt{p}[1] \to \mathtt{p}[i] : \left\{ \begin{array}{l} \mathsf{on} : \mathtt{p}[i] \to \mathtt{p}[4] : \mathsf{on}; \mathtt{p}[4] \to \mathtt{p}[6] : \mathsf{on}, \\ \mathsf{off} : \mathtt{p}[i] \to \mathtt{p}[5] : \mathsf{on}; \mathtt{p}[5] \to \mathtt{p}[6] : \mathsf{on} \end{array} \right\}; \\
\quad (\mathtt{p}[i] \to \mathtt{p}[1] : \langle \mathsf{unit} \rangle; \mathtt{p}[4] \to \mathtt{p}[i] : \langle \mathsf{unit} \rangle; \mathtt{p}[5] \to \mathtt{p}[i] : \langle \mathsf{unit} \rangle)\}; \\
\quad (\mathtt{p}[6] \to \mathtt{p}[4] : \langle \mathsf{unit} \rangle; \mathtt{p}[5] \to \mathtt{p}[4] : \langle \mathsf{unit} \rangle); \\
\quad \texttt{foreach}(n < m)\{\mathtt{p}[1] \to \mathtt{p}[6] : \langle \mathsf{real} \rangle\}; \mathtt{p}[6] \to \mathtt{p}[1] : \langle \mathsf{unit} \rangle; \\
\quad);\mathbf{x}
\end{array}
$$

The processes are similarly defined. For example, $\mathtt{P}[1]$ can be implemented as below (we assume c_{ij} denotes the confirmation channel from $\mathtt{P}[i]$ to $\mathtt{P}[j]$):

$$
\begin{array}{l}
\mu X.(\texttt{foreach}(i \in \{2,3\})\{\texttt{foreach}(n < m)\{s_{1i} \lhd l\}; c_{i1}?()\}; \\
\quad \texttt{foreach}(n < m)\{s_{16}!\langle V \rangle\}; c_{61}?()); X;
\end{array}
$$

Since each confirmation is sent independently, there are more parallel executions (i.e. independent loops) compared to the first one.