

Buffered Communication Analysis in Distributed Multiparty Sessions^{*}

Pierre-Malo Deniérou 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, 11, 15, 22], 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 [10, 20], can specify communication protocols by describing the sequences and types of read, write and choices on a given channel. For example, type $T_0 = !\langle \text{nat} \rangle; !\langle \text{string} \rangle; ?\langle \text{real} \rangle; \text{end}$, 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.

$T_1 = !\langle \text{nat} \rangle; ?\langle \text{real} \rangle; !\langle \text{string} \rangle; \text{end}$ 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 [7] and formally studied in [8] 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 = $\mu X.s_1?(x); s_2! \langle \text{Orange} \rangle; X$

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

Example (b)

(B₁) Bob₁ = $\mu X.s_1?(x); X$

(C₁) Carol₁ = $\mu X.s_3! \langle 2.4 \rangle; X$



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 than one value at any given time.

However, if we change Bob and Carol to Bob₁ and Carol₁ 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 Bob₁ and Carol₁.

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 $T_a = \mu x.s_1! \langle \text{nat} \rangle; s_3? \langle \text{real} \rangle; x$ (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, 11, 15, 22] for the analysis:

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

The above type represents the global interaction between Alice-Bob-Carol in (a) where Alice \rightarrow Bob: $s_1 \langle \text{nat} \rangle$; 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-Bob₁-Carol₁ in (b) lacks the second Bob \rightarrow 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)

(A₂) Alice₂ = $\mu X.s_1! \langle 1 \rangle; s_3?(x); s_4! \langle x+1 \rangle; X$


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

(C₂) Carol₂ = $\mu X.s_2?(x); s_3! \langle 2.4 \rangle; s_4?(y); X$

Example (d)

(A₃) Alice₃ = $\mu X.s_1! \langle 1 \rangle; s_3?(x); s_2! \langle x+1 \rangle; X$

(C₃) Carol₃ = $\mu X.s_2?(x); s_3! \langle 2.4 \rangle; s_2?(y); X$



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-Bob₁-Carol₁ 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:

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

The revised processes following G' are given as:

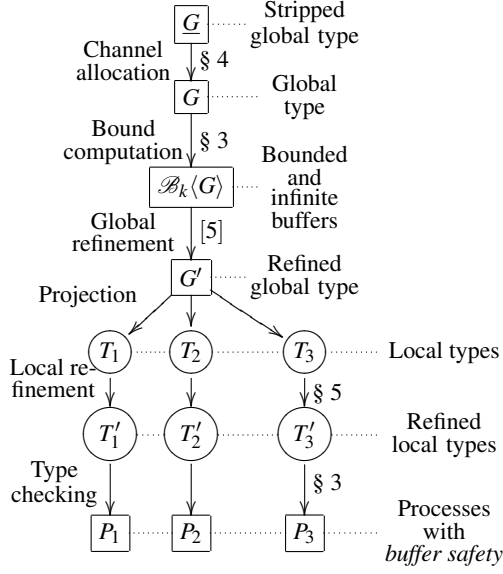
$$\text{Bob}_4 = \mu X.s_1?(x); s_1?(x); s_2! \langle \text{Signal} \rangle; X \quad \text{Carol}_4 = \mu X.s_3! \langle 2.4 \rangle; s_2?(x).s_3! \langle 2.4 \rangle; X$$

Local refinement for multiparty messaging optimisations. The last issue is about flexible local refinement (optimisations) based on [14, 15]. 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 .

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

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 $s_2! \langle \text{Orange} \rangle; s_2! \langle \text{Orange} \rangle; \dots; s_2! \langle \text{Orange} \rangle; \text{Bob}$. 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). Due to the space limitation, we omit the global refinement. A full version which includes the global refinement, omitted definitions, examples and proofs, and a prototype for computing the upper bound on the buffer size are available [5].



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 (long version [5]).
4. Once the global type has been projected to local types, local refinement can optimise the distributed execution of the participants' processes (§ 5).
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 (§ 5.1).

2 Asynchronous Multiparty Sessions

Syntax. We start from the π -calculus for multiparty sessions from [11] with unbounded and bounded buffers. Base sets and the grammars are given below.

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

$\bar{a}_{[2..n]}(\tilde{s}^m).P$ initiates, through a shared name a , a new session s_i with buffer size m_i ($1 \leq n \leq \infty$) with other participants, each of the form $a_{[p]}(\tilde{s}).Q$ with $1 \leq p \leq n-1$. 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, \dots (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 [11]. We often omit n from $s^n:\tilde{h}$, $\mathbf{0}$, and unimportant arguments e.g. $s!\langle\tilde{s}\rangle$ and $s?(\tilde{x});P$. An *initial* process does not contain any runtime syntax (buffers and session hiding).

Reductions. A selection of reduction rules is given below.

$$\begin{array}{l}
\bar{a}[2..n](\tilde{s}^n).P_1 \mid a[2].P_2 \mid \dots \mid a[n].P_n \rightarrow (v\tilde{s})(P_1 \mid P_2 \mid \dots \mid P_n \mid s_1^{n_1}:\emptyset \mid \dots \mid s_m^m:\emptyset) \\
s!(\tilde{e});P \mid s^n:\tilde{h} \rightarrow P \mid s^n:\tilde{h}\cdot\tilde{v} \quad (n \geq |\tilde{h}|, e_i \downarrow v_i) \quad s?(\tilde{x});P \mid s^n:\tilde{v}\cdot\tilde{h} \rightarrow P[\tilde{v}/\tilde{x}] \mid s^n:\tilde{h} \\
s!(\tilde{r});P \mid s^n:\tilde{h} \rightarrow P \mid s^n:\tilde{h}\cdot\tilde{r} \quad (n \geq |\tilde{h}|) \quad s?(\tilde{r});P \mid s^n:\tilde{r}\cdot\tilde{h} \rightarrow P \mid s^n:\tilde{h} \\
s \triangleleft l;P \mid s^n:\tilde{h} \rightarrow P \mid s^n:\tilde{h}\cdot l \quad (n \geq |\tilde{h}|) \quad s \triangleright \{l_i: P_i\}_{i \in I} \mid s^n:l_j\cdot\tilde{h} \rightarrow P_j \mid s^n:\tilde{h} \quad (j \in I)
\end{array}$$

The first rule describes the initiation of a new session among 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 [11]. The input rules perform the complementary operations. Processes are considered modulo a structural equivalence \equiv , whose definition is standard (e.g. $\mu X.P \equiv P[\mu X.P/X]$) [11].

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', \dots , describes the whole conversation scenario of a multiparty session as a type signature. Our starting syntax is from [11].

$$\begin{array}{ll}
G, G' ::= p \rightarrow p' : k \langle U \rangle ; G' & \text{values} \\
| p \rightarrow p' : k \{l_j : G_j\}_{j \in J} & \text{branching} \\
| \mu \mathbf{x}.G \mid \mathbf{x} \mid \text{end} & \text{recursion, end}
\end{array}
\quad
\begin{array}{ll}
U, U' ::= \tilde{S} \mid T @ p & \text{sorts, session} \\
S, S' ::= \text{bool} \mid \text{nat} \mid G & \text{base, shared}
\end{array}$$

Type $p \rightarrow p' : k \langle U \rangle ; G'$ says that participant p sends a message of type U on the channel k (represented as a natural number) so that participant 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* $T @ p$, representing a local type T assigned to participant 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 $p \rightarrow p' : k \{l_j : G_j\}_{j \in J}$ says that participant p can invoke one of the l_j labels on channel k (for participant p' to read) and that interactions described in G_j follow. We require $p \neq p'$ to prevent self-sent messages. Type $\mu \mathbf{x}.G$ is for recursive protocols, assuming the type variables $(\mathbf{x}, \mathbf{x}', \dots)$ 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 $\text{chans}(G)$ and $\text{prins}(G)$ 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 $p \rightarrow p' : k \langle U \rangle ; G'$ or $p \rightarrow p' : k \{l_j : G_j\}_{j \in J}$ with a node name (n_1, n_2, \dots) . Then, we inductively define a function node_G that gives a node n_k (or the special node end) for each of the syntactic subterm of G as follows:

$$\begin{aligned} \text{node}_G(\text{end}) &= \text{end} & \text{node}_G(n_i : p \rightarrow p' : k \langle U \rangle ; G') &= n_i \\ \text{node}_G(\mu \mathbf{x}.G') &= \text{node}_G(G') & \text{node}_G(n_j : p \rightarrow p' : k \{l_j : G_j\}_{j \in J}) &= n_j \\ \text{node}_G(\mathbf{x}) &= \text{node}_G(\mu \mathbf{x}.G') & & \text{(if the binder of } \mathbf{x} \text{ is } \mu \mathbf{x}.G' \in G) \end{aligned}$$

We define G as a session graph in the following way: for each subterm of G of the form $n : p \rightarrow p' : k \langle U \rangle ; G'$, we have an edge from n to $\text{node}_G(G')$, and for each subterm of G of the form $n' : p \rightarrow p' : k \{l_j : G_j\}_{j \in J}$, we have edges from n' to each of the $\text{node}_G(G_j)$ for $j \in J$. We also define the functions $\text{pfx}(n_i)$ and $\text{ch}(n_i)$ that respectively give the prefix $(p \rightarrow p' : k)$ and channel (k) that correspond to n_i . For a global type G , $\text{node}_G(G)$ distinguishes the *initial* node. $\text{size}(G)$ 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). n_1 is the initial node.



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 $n \prec n'$ (omitting G). Paths in this session graph are referred to by the sequence of nodes they pass through: a path $n_0 \prec \dots \prec n_n$ can be written more concisely $n_0 \dots n_n$ or \bar{n} when there is no ambiguity. We say that a path $n_0 \dots n_n$ has *suffix* $n_i \dots n_n$ 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_{IO} , defined below:

$$n_1 \prec_{\text{IO}} n_2 \text{ if } n_1 \prec\!\!\prec n_2 \text{ and } \text{pfx}(n_1) = p_1 \rightarrow p : k_1 \text{ and } \text{pfx}(n_2) = p \rightarrow p_2 : k_2 \text{ with } k_1 \neq k_2$$

The relation \prec_{IO} asserts the order between a reception by a principal and the next message it sends. An *input-output dependency* (IO-dependency) from n_1 to n_n is a chain $n_1 \prec_{\text{IO}} \dots \prec_{\text{IO}} n_n$ ($n \geq 1$).


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 $n \in G$ and for every cycle \tilde{n} for \prec such that $\text{ch}(n) = k$ and $n \in \tilde{n}$, there exists a cycle for \prec_{IO} that starts from n and only involves nodes from \tilde{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 $O(\text{size}(G)^3)$. The proof relies on the fact that establishing all IO-dependencies of a given session has $O(\text{size}(G)^3)$ 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 $n_1 n_2 n_3 n_4$ is part of the IO-cycles $n_1 n_2 n_3$ or $n_3 n_4$, 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 $\tilde{n} = n_0 \dots n_n n$ is a path in G , the property $\text{Reset}(\tilde{n})$ holds if there exist $0 \leq i_0 < \dots < i_j \leq n$ ($j \geq 1$) such that $n_{i_0} \prec_{\text{IO}} \dots \prec_{\text{IO}} n_{i_j} \prec_{\text{IO}} n$ and $\text{ch}(n_{i_0}) = \text{ch}(n)$. One practical instance of the nodes $\{n_{i_0}, \dots, n_{i_j}, n\}$ is called the reset nodes of \tilde{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 $\mathcal{B}_k(G)$ 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 $\mathcal{B}_k(G) = \mathcal{B}_k(0, \emptyset, \varepsilon, n_0)$ for n_0 the initial node of G .

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

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 $\mathcal{B}_k(m, \mathcal{P}, \tilde{n}, n)$) 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 \mathcal{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 \mathcal{P} for readability):

$\mathcal{B}_{s_2}(0, \varepsilon, n_1)$	max	explanation
$= \mathcal{B}_{s_2}(0, n_1, n_2)$	0	$s_1 \neq s_2$
$= \max(\mathcal{B}_{s_2}(1, n_1 n_2, n_3), \mathcal{B}_{s_2}(1, n_1 n_2, n_5))$	1	$\neg \text{Reset}(n_1 n_2), \neg \text{Reset}(n_1 n_3)$
$= \max(\mathcal{B}_{s_2}(1, n_1 n_2 n_3, n_4), \mathcal{B}_{s_2}(1, n_1 n_2 n_5, \text{end}))$	1	$s_3 \neq s_2$
$= \max(\mathcal{B}_{s_2}(1, n_4, n_1), 0)$	1	$\text{Reset}(n_1 n_2 n_3 n_4)$
$= \mathcal{B}_{s_2}(1, n_4 n_1, n_2)$	1	$s_1 \neq s_2$
$= \max(\mathcal{B}_{s_2}(2, n_4 n_1 n_2, n_3), \mathcal{B}_{s_2}(2, n_4 n_1 n_2, n_5))$	2	$\neg \text{Reset}(n_4 n_1 n_2), \neg \text{Reset}(n_4 n_1 n_3)$
$= \max(\mathcal{B}_{s_2}(2, n_4 n_1 n_2 n_3, n_4), \mathcal{B}_{s_2}(2, n_4 n_1 n_2 n_5, \text{end}))$	2	$s_3 \neq s_2$
$= \max(\mathcal{B}_{s_2}(1, n_4, n_1), 0)$	2	$\text{Reset}(n_4 n_1 n_2 n_3 n_4)$

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 $n_1 n_2$ 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 $\text{Reset}(n_1 n_2 n_3 n_4)$ holds (there is $n_2 \prec_{\text{IO}} n_3 \prec_{\text{IO}} n_4$), 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 its successor n_2 does. We thus check the reset property on the path $n_4 n_1 n_2$, 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 $n_4 n_1 n_2 n_3$. The next prefix, n_4 , uses s_2 and $\text{Reset}(n_4 n_1 n_2 n_3 n_4)$ 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, \dots, P_n can be type-checked against T_1, \dots, T_n , they are automatically guaranteed to interact properly, without communication mismatch (communication safety) nor getting stuck inside a session (progress) [11]. 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.

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

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 $k \oplus \{l_i : T_i\}_{i \in I}$ 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 $k \& \{l_i : T_i\}_{i \in I}$ 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 $G \upharpoonright p$ (called *projection of G onto p*) and defined in [11, 22]. For example, $(p \rightarrow p' : k \langle U \rangle; G') \upharpoonright p = k! \langle U \rangle; (G' \upharpoonright p)$, $(p \rightarrow p' : k \langle U \rangle; G') \upharpoonright p' = k? \langle U \rangle; (G' \upharpoonright p')$ and $(p \rightarrow p' : k \langle U \rangle; G') \upharpoonright q = (G' \upharpoonright q)$. We take an *equi-recursive* view, not distinguishing between $\mu x.T$ and its unfolding $T[\mu x.T/x]$.

Linearity. To avoid race conditions and conflicts between typed processes, we build on the definition of linearity from [11]. The relations \prec_{II} and \prec_{OO} are defined by:

$$\begin{aligned} n_1 \prec_{\text{II}} n_2 & \text{ if } n_1 \ll n_2 \text{ and } \text{pfx}(n_1) = p_1 \rightarrow p : k_1 \text{ and } \text{pfx}(n_2) = p_2 \rightarrow p : k_2 \text{ s.t. } p_1 \neq p_2 \Leftrightarrow k_1 \neq k_2 \\ n_1 \prec_{\text{OO}} n_2 & \text{ if } n_1 \ll n_2 \text{ and } \text{pfx}(n_1) = p \rightarrow p_1 : k_1 \text{ and } \text{pfx}(n_2) = p \rightarrow p_2 : k_2 \text{ s.t. } p_1 \neq p_2 \Rightarrow k_1 \neq k_2 \end{aligned}$$

The three relations \prec_{IO} , \prec_{II} and \prec_{OO} are used to characterise the authorised sequences of actions. An *input dependency (I-dependency)* from n_1 to n_2 is a chain $n_1 \prec_{\phi_1} \cdots \prec_{\phi_n} n_2$ ($n \geq 1$) such that $\phi_i = \text{IO}$ for $1 \leq i \leq n-1$ and $\phi_n = \text{II}$. An *output dependency (O-dependency)* from n_1 to n_2 is a chain $n_1 \prec_{\phi_1} \cdots \prec_{\phi_n} n_2$ ($n \geq 1$) such that $\phi_i \in \{\text{OO}, \text{IO}\}$. These dependency relations are respectively written \ll_{II} and \ll_{OO} . G is *linear* (written $\text{Lin}(G)$) if, whenever two nodes $n_1 \ll n_2$ use the same channel k , the dependencies $n_1 \ll_{\text{II}} n_2$ and $n_1 \ll_{\text{OO}} n_2$ hold. If G carries other global types, we inductively demand the same. Examples can be found in [5, 11]. 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: $\Gamma ::= \emptyset \mid \Gamma, u : S \mid \Gamma, X : \Delta$ and $\Delta ::= \emptyset \mid \Delta, \tilde{s}^m : \{T @ p\}_{p \in I}$. A *sorting* (Γ, Γ', \dots) is a finite map from names to sorts and from process variables to sequences of sorts and types. *Typing* (Δ, Δ', \dots) 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 $\mathcal{B}_{s_i} \langle G \rangle$ calculated from G .

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

The type for \tilde{s} is the *first* projection of the declared global type for a in Γ . The end-point type $(G \upharpoonright p) @ p$ means that the participant p has $G \upharpoonright p$, which is the projection of G onto p , as its end-point type. The condition $\mid \tilde{s} \mid = \text{chans}(G)$ means the number of session channels meets those in G . The condition $\mathcal{B}_k \langle G \rangle = m_k$ 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 [11]. Note that since $\mathcal{B}_k \langle G \rangle$ is decidable, type-checking for processes with type annotations is decidable [11, 22].

The rest of the typing system for programs and one for runtime are similar with those in [11] ([5]). 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 $k! \langle U \rangle; T @_{\mathbf{p}}, k? \langle U \rangle; T' @_{\mathbf{q}} \xrightarrow{k} T @_{\mathbf{p}}, T' @_{\mathbf{q}}$. The key lemma about the correspondence between buffer size and reduction follows. We set $\llbracket G \rrbracket$ to be the family $\{(G \upharpoonright_{\mathbf{p}}) @_{\mathbf{p}} \mid \mathbf{p} \in G\}$. Regarding each type in $\llbracket G \rrbracket$ as the corresponding regular tree, we can define $\prec, \prec_{\text{II}}, \prec_{\text{IO}}$ and \prec_{OO} among its prefixes precisely as we have done for G .

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

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

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 $\mathcal{B}_k \langle G \rangle \geq \mathcal{B}_k \langle G' \rangle$ with $\Delta(\tilde{s}) = \llbracket G \rrbracket$, $\Delta'(\tilde{s}) = \llbracket G' \rrbracket$ 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:

$$\begin{aligned} n \leq |\tilde{h}| &\Rightarrow s! \langle \tilde{e} \rangle; P \mid s^n : \tilde{h} \rightarrow_{\text{Err}}, s! \langle \tilde{e} \rangle; P \mid s^n : \tilde{h} \rightarrow_{\text{Err}}, s \triangleleft l; P \mid s^n : \tilde{h} \rightarrow_{\text{Err}} \\ P \rightarrow_{\text{Err}} &\Rightarrow P \mid Q \rightarrow_{\text{Err}}, (v a)P \rightarrow_{\text{Err}}, (v \tilde{s})P \rightarrow_{\text{Err}}, P \equiv Q \rightarrow_{\text{Err}} \end{aligned}$$

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

4 Channel Attribution

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

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

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 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: $\forall n, n' \in G$, $\text{ch}(n) = \text{ch}(n') \iff n = n'$. 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 , $\mathcal{B}_k \langle G \rangle = 1$; (3) for the finite channels k of G , $\sum_k \mathcal{B}_k \langle G \rangle \leq \text{size}(G)$.*

Channel equalities. As well as values of the ch function, allocations can be seen as partitions of the set of nodes $\{n\}_{n \in G}$. 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 $\mathcal{B}_k \langle G \rangle = \infty$ then $\forall E, \mathcal{B}_k \langle G_E \rangle = \infty$.

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 [19] 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, 8] follow this allocation. Formally, the principal allocation is such that: $\forall n, n' \in G$ s.t. $\text{pfx}(n) = \text{p} \rightarrow \text{q} : k$ and $\text{pfx}(n') = \text{p}' \rightarrow \text{q}' : k'$, $(k = k' \iff \text{p} = \text{p}' \wedge \text{q} = \text{q}')$.

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

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 n_0 , and a successor function succ over the nodes, the function $\mathcal{A}_0^{\text{K}}(n_0)$ is defined by:

$$\mathcal{A}_E^{\text{K}}(n) = \mathcal{A}_{E'}^{\text{K}'}(n') \text{ where } \begin{cases} \text{succ}(n) = n' \\ \text{ch}(n) = k \quad \wedge E' = \begin{cases} E \cup \{k = k'\} & \text{if } \exists k' \in \text{K}, \text{Lin}(G_{E \cup \{k = k'\}}) \\ E & \text{otherwise} \end{cases} \\ \text{K}' = \text{K} \cup \{k\} \end{cases}$$

$\mathcal{A}_E^{\text{K}}(\text{end}) = E$

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

The early and late allocations are not optimal in terms of total memory requirements (computed by $\sum_k \mathcal{B}_k \langle G \rangle$ 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.
$n_0 : A \rightarrow B;$	k_0	k_0	k_0	k_0
$n_1 : B \rightarrow A;$	k_1	k_1	k_1	k_1
$n_2 : A \rightarrow B;$	k_2	k_0	k_0	k_0
$n_3 : A \rightarrow B;$	k_3	k_0	k_0	k_1
$n_4 : A \rightarrow C;$	k_4	k_2	k_2	k_2
$n_5 : C \rightarrow B;$	k_5	k_3	k_1	k_3
$n_6 : B \rightarrow C;$	k_6	k_4	k_2	k_1
$n_7 : B \rightarrow C$	k_7	k_4	k_2	k_2
Nb channels	8	5	3	4
Memory Req.	8	7	5	5

5 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. Our recent work [14, 15] 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 Bob₅ 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 $\text{dep}(k)$ 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 $\{N_0, \dots, N_n\}$ of the set of nodes of G . This partition should satisfy the two properties: $\forall n \in N_i, n' \in N_j, \text{ch}(n) = \text{ch}(n') \Rightarrow N_i = N_j$ and $\forall n \in N_i, \text{dep}(\text{ch}(n)) \subset N_i$. 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 $N_1 = \{n_1, n_2\}$ and $N_2 = \{n_3, n_4\}$: 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.

$$\begin{array}{c}
 n_1 : \text{Alice} \rightarrow \text{Bob} : s_1^{N_1} \\
 \vdots \\
 n_4 : \text{Bob} \rightarrow \text{Alice} : s_4^{N_2}
 \end{array}
 \begin{array}{c}
 \xrightarrow{\text{IO}} \\
 \xrightarrow{\text{IO}} \\
 \xrightarrow{\text{IO}} \\
 \xrightarrow{\text{IO}}
 \end{array}
 \begin{array}{c}
 n_2 : \text{Bob} \rightarrow \text{Alice} : s_2^{N_1} \\
 \vdots \\
 n_3 : \text{Alice} \rightarrow \text{Bob} : s_3^{N_2}
 \end{array}
 \left| \begin{array}{l}
 T_{\text{Alice}} = \mu \mathbf{x}. s_1^{N_1} !; s_2^{N_1} ?; s_3^{N_2} !; s_4^{N_2} ?; \mathbf{x} \\
 T_{\text{Bob}} = \mu \mathbf{x}. s_1^{N_1} ?; s_2^{N_1} !; s_3^{N_2} ?; s_4^{N_2} !; \mathbf{x} \\
 T_{\text{Alice}}^{\text{opt}} = \mu \mathbf{x}. s_1^{N_1} !; s_3^{N_2} !; s_2^{N_1} ?; s_4^{N_2} ?; \mathbf{x}
 \end{array} \right.$$

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 asyn-

chronous than (or more optimised than) T' . The main rule is:

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

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 $T_1 \leq_c T_2$ as a form of type simulation, following [14, 15]. The important fact is that \leq_c does not alter buffer sizes: suppose $\llbracket G \rrbracket = \{T @ p\}_p$ with $T @ p = (G \upharpoonright p) @ p$ and $p \in G$. Assume $T @ p \leq_c T' @ p$ with $\llbracket G' \rrbracket = \{T' @ p\}_p$. Then $\mathcal{B}_k \langle G \rangle = \mathcal{B}_k \langle G' \rangle$. 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 $s_3^{N_2} !$ ($T_{\text{Alice}}^{opt} \leq_c T_{\text{Alice}}$) since $N_1 \cap N_2 = \emptyset$, keeping the size of each buffer one. Hence process typable by T_{Alice}^{opt} can safely send message at s_3 before input at s_2 . In Alice-Bob₅-Carol from § 1, the original global type G annotated by IO-chains has only one partition $N = \{n_1, n_2, n_3\}$:

$$\mu x. \text{Alice} \rightarrow \text{Bob}: s_1^N \langle \text{nat} \rangle; \text{Bob} \rightarrow \text{Carol}: s_2^N \langle \text{string} \rangle; \text{Carol} \rightarrow \text{Alice}: s_3^N \langle \text{real} \rangle; x$$

Bob's local type is $\mu x. s_1^N ? \langle \text{nat} \rangle; s_2^N ! \langle \text{string} \rangle; x$, which prevents any optimisation \ll by (OI). Hence, Bob₅ is not typable. Some typable examples are given in the next section.

5.1 Application: Multi-buffering algorithm

The double buffering algorithm [6] is widely used in high-performance and multicore computing. We generalise this algorithm to *multi-buffering* [16], and solve an open issue in our previous work [15, § 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 [22] (i.e. where $\text{foreach}(i < n)\{P[i]\}$ means we iterate $P[i]$ for $0 \leq i < n$):

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

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

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

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, \dots, 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 5.1 (*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 $r_0!\langle \rangle; \dots; \text{foreach}(i < n)\{r_i!\langle \rangle; s_i?(x_i); t_i?; u_i!\langle x_i \rangle\}$ (which is not typable in our system), then all buffers are forced to hold 2 units. This unsafe optimisation is typable in [15] but prevented here by Proposition 5.1. In [5], we also deal with a use case of MPSoC buffer allocation from [4], with branching and iterations [22], and verify it by applying all of the previously described methods.

6 Related Work

Checking buffer bounds based on global specifications has been studied through Petri nets and Synchronous data flow. Recent advances [9] in the study of Kahn Process Networks (KPN) have improved Parks’s algorithm [18] 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 [12, 22] 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 [17]. Gay & Vasconcelos [8] 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 [21] describe a polynomial method to infer buffer bounds of a concurrent language through program analysis using linear programming techniques, improving on previous work in [13], see [21, § 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 [12]). 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 [21, § 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 [8, 21] 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 $\llbracket G \rrbracket$ [15], Subject Reduction Theorem can be

proved using our invariance method noting that we use $\llbracket G \rrbracket$ 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. P.-M. Deniérou and N. Yoshida. Buffered communication analysis in distributed multiparty sessions. Full version, Prototype at <http://www.doc.ic.ac.uk/~pmalo/multianalysis>.
6. 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.
7. 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.
8. S. Gay and V. T. Vasconcelos. Linear type theory for asynchronous session types. *JFP*, 2009.
9. M. Geilen and T. Basten. Requirements on the Execution of Kahn Process Networks. In *ESOP*, volume 2618 of *LNCS*, pages 319–334. Springer, 2003.
10. 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.
11. K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL*, pages 273–284, 2008.
12. R. Hu, N. Yoshida, and K. Honda. Session-Based Distributed Programming in Java. In *ECOOP’08*, volume 5142 of *LNCS*, pages 516–541, 2008.
13. N. Kobayashi, M. Nakade, and A. Yonezawa. Static analysis of communication for asynchronous concurrent programming languages. In *SAS*, *LNCS* 983, pages 225–242, 1995.
14. 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.
15. 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.
16. Multi-buffering. http://en.wikipedia.org/wiki/Multiple_buffering.
17. H. Nielson and F. Nielson. Higher-order concurrent programs with finite communication topology. In *POPL*, pages 84–97, 1994.
18. T. Parks. *Bounded Scheduling of Process Networks*. PhD thesis, California Berkeley, 1995.
19. G.-C. Rota. The number of partitions of a set. *Amer. Math. Monthly*, 71:498–504, 1964.
20. 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.
21. T. Terauchi and A. Megacz. Inferring channel buffer bounds via linear programming. In *ESOP*, volume 4960 of *LNCS*, pages 284–298. Springer, 2008.
22. N. Yoshida, P.-M. Deniérou, A. Bejleri, and R. Hu. Parameterised multiparty session types. In *FoSSaCs*, volume 6014 of *LNCS*, pages 128–145, 2010.