

Global Principal Typing in Partially Commutative Asynchronous Sessions

Dimitris Mostrous¹, Nobuko Yoshida¹, and Kohei Honda²

¹ Department of Computing, Imperial College London

² Department of Computer Science, Queen Mary, University of London

Abstract. We generalise a theory of multiparty session types for the π -calculus through asynchronous communication subtyping, which allows partial commutativity of actions with maximal flexibility and safe optimisation in message choreography. A sound and complete algorithm for the subtyping relation, which can calculate conformance of optimised end-point processes to an agreed global specification, is presented. As a complementing result, we show a type inference algorithm for deriving the principal global specification from end-point processes which is minimal with respect to subtyping. The resulting theory allows a programmer to choose between a top-down and a bottom-up style of communication programming, ensuring the same desirable properties of typable processes.

1 Introduction

Programs which communicate by asynchronous message passing are abundant in critical computing scenes, from a simple web-service application between two parties to a global financial network hosting thousands of nodes and billions of messages per year. The design of such programs, which may be developed in geographically disparate sites, demands a clear high-level specification of their conversation structure, against which participating programs can be validated (*conformance*). Further such specifications may change during development (*refinement*), and might even need to be synthesised from individual endpoint programs, against which updated programs can be further validated (*synthesis of global specifications*).

This paper develops a new theory of multiparty session types [1, 2, 4, 5, 13, 23], which can handle uniformly these three concerns by seamlessly integrating the *top-down* and *bottom-up* strategies for the development of communication-centred software. The methodology for distributed programming put forward in [1, 13] centres on the concept of a *global type* which plays the role of type signature for distributed communications, presenting an abstract high-level description of the protocol that all the participants have to honour when an actual conversation takes place. Building on this framework, we propose the following two strategies for communication programming:

Top-Down Approach: Once this signature G is agreed upon by all parties as a global protocol, a local protocol from each party's viewpoint (*local type* T_i) is generated as a projection of G to each party. Then each local type T_i can be *locally refined* to, say, T'_i , possibly giving a more *optimised* protocol, realised as a program, say, P_i . If all the resulting local programs P_1, \dots, P_n can be type-checked against refined T'_1, \dots, T'_n , then they

are automatically guaranteed to interact properly, without communication mismatch or getting stuck inside a session, precisely following the intended scenario.

Bottom-Up Approach: In this case the programmers may work based on an informal understanding of shared conversation structures, which, after appropriate development, will get reified into a formal global protocol by synthesis of local behaviours of all endpoint programs: first, a type T_i is inferred from each program P_i , then a new global specification is synthesised from T_1, \dots, T_n . If this specification is validated to satisfy certain conditions, P_1, \dots, P_n are guaranteed to interact properly. This process can be repeated incrementally, using a succession of synthesised types as *globally refined* protocols.

This paper presents a general and rigorous foundation for these two approaches and their seamless integration, based on multiparty session types. For the automatic refinement, we introduce *asynchronous communication subtyping* over local types, which allows permutation of actions to increase efficiency, while ensuring type-soundness and communication-safety. As an example, suppose we are using an asynchronous communication transport where the message order is preserved but the sending is non-blocking, as in TCP. Let us assume the following three simple processes:

$$P_1 \stackrel{\text{def}}{=} t?(y_1);s!\langle 5 \rangle;s!\langle \text{apple} \rangle;Q_1 \quad P_2 \stackrel{\text{def}}{=} b?(y_2);t!\langle 7 \rangle;Q_2 \quad P_3 \stackrel{\text{def}}{=} s?(z_1);s?(z_2);Q_3$$

where $s?(y)$ is an input and $s!\langle 5 \rangle$ is an output via channel s ; and “;” is sequential composition. Then first P_2 gets the value at b ; then P_2 sends 7 to P_1 ; finally P_1 sends 5 and “apple” to P_3 preserving the order. We note that P_3 is blocked until b is fired at P_2 . To execute P_3 ahead, P_1 might be locally optimised since y_1 does not bind the subsequent outputs at s . We can similarly optimise P_2 . The resulting processes given below still preserve linearity and proper communication structures.

$$P'_1 \stackrel{\text{def}}{=} s!\langle 5 \rangle;s!\langle \text{apple} \rangle;t?(y_1);Q_1 \quad P'_2 \stackrel{\text{def}}{=} t!\langle 7 \rangle;b?(y_2);Q_2 \quad P_3 \stackrel{\text{def}}{=} s?(z_1);s?(z_2);Q_3$$

Asynchronous communication subtyping specifies safe permutations of actions, by which we can refine a local protocol to maximise asynchrony without violating the global protocol. For example, in the above case, P'_1 is given local type $s!\langle \text{nat} \rangle;s!\langle \text{string} \rangle;t?\langle \text{nat} \rangle;T$ which is a subtype of $t?\langle \text{nat} \rangle;s!\langle \text{nat} \rangle;s!\langle \text{string} \rangle;T$ projected from the global type. Hence optimisations can be checked locally. The idea of this subtyping is intuitive, but it requires delicate formal formulations due to the presence of recursive types and branching/selection session types, whose combinations are vital for typing many practical protocols [12, 22]. This subtlety is because type-permutations *affect the structures of session types*, which makes straightforward constructions following the preceding literature [10, 20] inapplicable. Intuitively, because partial commutativity is defined between *a sequence of actions*, it may require *more than one unfolding* of recursive types to find a match. However, this calculation can be made automatic by an *algorithmic subtyping* which completely characterises the semantic notion of subtyping and can be used to effectively (in)validate conformance of an optimised local type to a global type.

For the bottom-up strategy, we formulate *principal global typing* by which we can synthesise the most general global type from untyped endpoint programs, or can check they can have no global type, i.e. their protocols are incompatible. The framework

uses graph-shaped global types which generalise the original syntactic global types, extending typability. Asynchronous communication subtyping plays a central role in the synthesis process. We demonstrate the use of the theory for the two strategies by providing correctness arguments for the development of a distributed parallel algorithm. A full version, containing more examples and detailed proofs, is available from <http://www.doc.ic.ac.uk/~mostrous/asynsub>.

2 Asynchronous Multiparty Sessions

Syntax. We use the π -calculus for multiparty sessions from [13], omitting polyadicity and delegation for simplicity. We use the following base sets: *shared names* or *names*, written a, b, x, y, z, \dots ; *session channels* or *channels*, written s, t, \dots ; *labels* (functioning like labels in labelled records), written l, l', \dots ; and *process variables*, written X, Y, \dots . For hiding, we use n for either a single shared name or a vector of session channels. Then *processes* (P, Q, \dots) and *expressions* (e, e', \dots) are given below:

$P ::= \bar{a}_{[2..n]}(\tilde{s}).P$	request		$P \mid Q$	parallel
$a_{[p]}(\tilde{s}).P$	acceptance		$\mathbf{0}$	inaction
$s!(e);P$	sending		$(\nu n)P$	hiding
$s?(x);P$	reception		$\text{def } D \text{ in } P$	recursion
$s \triangleleft l;P$	selection		$X \langle \tilde{e}\tilde{s} \rangle$	process call
$s \triangleright \{l_i : P_i\}_{i \in I}$	branching		$s:\tilde{h}$	message queue
$\text{if } e \text{ then } P \text{ else } Q$	conditional			
$e ::= v \mid e \text{ and } e' \mid \text{not } e \dots$	expressions	$h ::= l \mid v$		message values
$v ::= a \mid \text{true} \mid \text{false} \dots$	values	$D ::= \{X_i(\tilde{x}_i\tilde{s}_i) = P_i\}_{i \in I}$		declaration

$\bar{a}_{[2..n]}(\tilde{s}).P$ initiates, through a shared name a , a new session with other participants, each of shape $a_{[p]}(\tilde{s}).Q$ with $2 \leq p \leq n$. The (bound) s_i in vector \tilde{s} are session channels 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; and by selection and branching (the former chooses one of the branches offered by the latter). $s:\tilde{h}$ is a *message queue* representing ordered messages in transit \tilde{h} with destination s (which may be considered as a network pipe in a TCP-like transport). The rest of the syntax is standard from [13]. We often omit $\mathbf{0}$, and unimportant arguments of sending/receiving, e.g. $s!(\)$ and $s?(\);P$.

Operational semantics. Some selected rules of reduction $P \rightarrow P'$ are given below:

$$\begin{aligned}
&\bar{a}_{[2..n]}(\tilde{s}).P_1 \mid a_{[2]}(\tilde{s}).P_2 \mid \dots \mid a_{[n]}(\tilde{s}).P_n \rightarrow (\nu \tilde{s})(P_1 \mid P_2 \mid \dots \mid P_n \mid s_1:\mathbf{0} \mid \dots \mid s_m:\mathbf{0}) \\
&s!(e);P \mid s:\tilde{h} \rightarrow P \mid s:\tilde{h} \cdot v \quad (e \downarrow v) \quad s?(x);P \mid s:\nu \cdot \tilde{h} \rightarrow P[v/x] \mid s:\tilde{h} \\
&s \triangleleft l;P \mid s:\tilde{h} \rightarrow P \mid s:\tilde{h} \cdot l \quad s \triangleright \{l_i : P_i\}_{i \in I} \mid s:l_j \cdot \tilde{h} \rightarrow P_j \mid s:\tilde{h} \quad (j \in I)
\end{aligned}$$

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 the private m fresh session channels s_i and the associated m empty queues ($\mathbf{0}$ denotes the empty string). The output rules enqueue a value and a label, respectively ($e \downarrow v$ denotes the evaluation of e to v). The input rules perform the complementary operations. Processes are considered modulo structural equivalence, \equiv , defined by the standard rules [13].

Global types. A *global type*, written G, G', \dots , describes the whole conversation scenario of a multiparty session as a type signature [13].

Global G ::=	$p \rightarrow p' : k \langle U \rangle ; G'$	values	$\mu \mathbf{t}.G$	recursive
	$p \rightarrow p' : k \{l_j : G_j\}_{j \in J}$	branching	\mathbf{t}	variable
	G, G'	parallel		end

Value U ::= $\text{bool} \mid \text{nat} \mid \dots \mid G$

Type $p \rightarrow p' : k \langle U \rangle ; G'$ says that participant p sends a message of type U to channel k (represented as a finite natural number) received by participant p' and then interactions described in G' take place. *Value types* range over U , and are either global types for shared names, or base values. Type $p \rightarrow p' : k \{l_j : G_j\}_{j \in J}$ says that participant p invokes one of the l_j labels on channel k (at participant p'), then interactions described in G_j take place. Type $\mu \mathbf{t}.G$ is for recursive protocols, assuming type variables $(\mathbf{t}, \mathbf{t}', \dots)$ are guarded in the standard way, i.e. they only occur under values or branches. We assume G in value types is closed, i.e. without free type variables. Type end represents the termination of a session. We often omit end and identify “ G, end ” and “ end, G ” with G . We stipulate that each channel can only be used among two parties (but maybe repeatedly), one party using it for input/branching while the other party for output/selection.¹

Local types. Local session types type-abstract sessions from each endpoint’s view.

Local T ::=	$k! \langle U \rangle ; T$	send	$k \& \{l_i : T_i\}_{i \in I}$	branching
	$k? \langle U \rangle ; T$	receive	$\mu \mathbf{t}.T \mid \mathbf{t}$	recursion
	$k \oplus \{l_i : T_i\}_{i \in I}$	selection		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 input. 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 types are standard. We say a type is *guarded* if it is neither a recursive type nor a type variable. (An occurrence of) a type constructor *not* under a recursive prefix in a recursive type is called *top-level action* (for example, $k_1! \langle U_1 \rangle$ and $k_2! \langle U_2 \rangle$ in $k_1! \langle U_1 \rangle ; k_2! \langle U_2 \rangle ; \mu \mathbf{t}.k_3! \langle U_3 \rangle ; \mathbf{t}$ are top-level, but $k_3! \langle U_3 \rangle$ in the same type is not). k is *the head of T* if k appears at the left-most occurrence of the top-level actions in T (e.g. $k_1! \langle U_1 \rangle$ is the head of the above type). The relation between global and local types is formalised by *projection*, written $G \upharpoonright p$ (called *projection of G onto p*), defined as in [13]. 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 write *Type* for the collection of all closed local types.

3 Asynchronous Partially Commutative Sessions

3.1 Asynchronous Communication Subtyping: Top-Level Actions

This section introduces and studies a basic theory of asynchronous session subtyping. Figure 1 defines the axioms for partial permutation of top-level actions for closed types,

¹ This condition dispenses with the need for linearity-check to ensure well-formedness [1].

$$\begin{array}{ll}
\text{(OI)} & k!\langle U \rangle; k'?\langle U' \rangle; T \ll k'?\langle U' \rangle; k!\langle U \rangle; T \\
\text{(OB)} & k!\langle U \rangle; k' \& \{l_j; T_j\}_{j \in J} \ll k' \& \{l_j; k!\langle U \rangle; T_j\}_{j \in J} \\
\text{(SI)} & k \oplus \{l_j; k'?\langle U \rangle; T_j\}_{j \in J} \ll k'?\langle U \rangle; k \oplus \{l_j; T_j\}_{j \in J} \\
\text{(SB)} & k \oplus \{l_i; k' \& \{l'_j; T_{ij}\}_{j \in J}\}_{i \in I} \ll k' \& \{l'_j; k \oplus \{l_i; T_{ij}\}_{i \in I}\}_{j \in J} \\
\text{(OO)} & k!\langle U \rangle; k'!\langle U' \rangle; T \ll k'!\langle U' \rangle; k!\langle U \rangle; T \\
\text{(II)} & k?\langle U \rangle; k'?\langle U' \rangle; T \ll k'?\langle U' \rangle; k?\langle U \rangle; T \\
\text{(SO)} & k \oplus \{l_i; k'!\langle U \rangle; T_i\}_{i \in I} \ll k'!\langle U \rangle; k \oplus \{l_i; T_i\}_{i \in I} \\
\text{(OS)} & k'!\langle U \rangle; k \oplus \{l_i; T_i\}_{i \in I} \ll k \oplus \{l_i; k'!\langle U \rangle; T_i\}_{i \in I} \\
\text{(SS)} & k \oplus \{l_i; k' \oplus \{l'_j; T_{ij}\}_{j \in J}\}_{i \in I} \ll k' \oplus \{l'_j; k \oplus \{l_i; T_{ij}\}_{i \in I}\}_{j \in J} \\
\text{(Tr)} & \frac{T_1 \ll T_2 \quad T_2 \ll T_3}{T_1 \ll T_3} \quad \text{(CO)} \frac{T \ll T'}{k!\langle U \rangle; T \ll k!\langle U \rangle; T'} \quad \text{(CI)} \frac{T \ll T'}{k'?\langle U \rangle; T \ll k'?\langle U \rangle; T'} \\
\text{(CB)} & \frac{\forall i \in I. T_i \ll T'_i}{k \& \{l_i; T_i\}_{i \in I} \ll k \& \{l_i; T'_i\}_{i \in I}} \quad \text{(CS)} \frac{\forall i \in I. T_i \ll T'_i}{k \oplus \{l_i; T_i\}_{i \in I} \ll k \oplus \{l_i; T'_i\}_{i \in I}} \quad \text{(E)} \text{end} \ll \text{end} \\
\text{(M)} & \mu t. T \ll \mu t. T
\end{array}$$

Fig. 1. Action Asynchronous Subtyping Rules ((BI, IB, BB) are omitted)

denoted \ll . We assume $k \neq k'$ for all the axioms. $T \ll T'$ is read: T is an *action-asynchronous subtype* of T' , and means T is more asynchronous than (or more optimised than) T' . We write $T \gg T'$ for $T' \ll T$.

A partial permutation is applied only to finite parts of the top-level actions (*without* unfolding recursive types); see Proposition 7. Note that we *cannot* exchange an input and output in the reverse direction of (OI) even for different channels. Consider: $P = s?(); r!\langle \rangle$ and $Q = s!\langle \rangle; r?()$. These processes interact correctly. If we permute the output and input of Q , we get $Q' = r?(); s!\langle \rangle$. Then the parallel composition ($P \mid Q'$) causes deadlock, losing progress. For the same reason, the reverse direction of (OB, SI, SB) is not allowed. By combining these input and output permutation rules, we can achieve a flexible local refinement for communications. For example, suppose $R = s?(x); r?(y); t!\langle 1 \rangle; t'!\langle y \rangle$ typed by $T_R = s?\langle \text{file} \rangle; r?\langle \text{bool} \rangle; t!\langle \text{nat} \rangle; t'!\langle \text{bool} \rangle; \text{end}$. We might wish to receive the (small) value via r first, and immediately forward to t' , then receive the (larger) file at s in the end: we can obtain $S = r?(y); t'!\langle y \rangle; t!\langle 1 \rangle; s?(x)$ typed by $T_S = r?\langle \text{bool} \rangle; t'!\langle \text{bool} \rangle; t!\langle \text{nat} \rangle; s?\langle \text{file} \rangle; \text{end}$, transformed from T_R (i.e. $T_S \ll T_R$) by using a combination of (OO, OI, II).

3.2 Asynchronous Communication Subtyping: Recursive Types

For handling recursive types in asynchronous subtyping, we extend the coinductive method in [20, § 2.3] and [10, § 3.3]. In particular, we need to modify the unfolding function for recursive types since \ll might be applicable to a *sequence of types* after unfolding of recursions under guarded prefixes. The resulting definition integrates \ll with the traditional session subtyping [10, 13]. For any recursive type T , $\text{unfold}^n(T)$ is the result of inductively unfolding the first recursion (even under guarded types) up to a fixed level of nesting.

Definition 1 (*n-time unfolding*)

$$\begin{aligned}
\text{unfold}^0(T) &= T \text{ for all } T & \text{unfold}^{1+n}(T) &= \text{unfold}^1(\text{unfold}^n(T)) \\
\text{unfold}^1(k!\langle U \rangle; T) &= k!\langle U \rangle; \text{unfold}^1(T) & \text{unfold}^1(k \oplus \{l_i : T_i\}_{i \in I}) &= k \oplus \{l_i : \text{unfold}^1(T_i)\}_{i \in I} \\
\text{unfold}^1(k?\langle U \rangle; T) &= k?\langle U \rangle; \text{unfold}^1(T) & \text{unfold}^1(k \& \{l_i : T_i\}_{i \in I}) &= k \& \{l_i : \text{unfold}^1(T_i)\}_{i \in I} \\
\text{unfold}^1(\mu \mathbf{t}. T) &= T[\mu \mathbf{t}. T / \mathbf{t}] & \text{unfold}^1(\mathbf{t}) &= \mathbf{t} & \text{unfold}^1(\text{end}) &= \text{end}
\end{aligned}$$

We also use $\text{unfold}^n(U)$ which is defined as $\text{unfold}^n(T)$ above.²

For example, $\text{unfold}^2(k?\langle U \rangle; \mu \mathbf{t}. k'!\langle U' \rangle; \mathbf{t}) = k?\langle U \rangle; k'!\langle U' \rangle; k'!\langle U' \rangle; \mu \mathbf{t}. k'!\langle U' \rangle; \mathbf{t}$. Note that, because our recursive types are contractive, $\text{unfold}^n(T)$ terminates. We can now introduce the central notion of asynchronous communication subtyping.

Definition 2. A relation $\mathfrak{R} \in \text{Type} \times \text{Type}$ is an asynchronous type simulation if $(T_1, T_2) \in \mathfrak{R}$ implies the following conditions:

- If $T_1 = \text{end}$, then $\text{unfold}^n(T_2) = \text{end}$.
- If $T_1 = k!\langle U_1 \rangle; T'_1$, then $\text{unfold}^n(T_2) \gg k!\langle U_2 \rangle; T'_2$, $(T'_1, T'_2) \in \mathfrak{R}$ and $(U_1, U_2) \in \mathfrak{R}$.
- If $T_1 = k?\langle U_1 \rangle; T'_1$, then $\text{unfold}^n(T_2) \gg k?\langle U_2 \rangle; T'_2$, $(T'_1, T'_2) \in \mathfrak{R}$ and $(U_2, U_1) \in \mathfrak{R}$.
- If $T_1 = k \oplus \{l_i : T_{1i}\}_{i \in I}$, then $\text{unfold}^n(T_2) \gg k \oplus \{l_j : T_{2j}\}_{j \in J}$ 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 $\text{unfold}^n(T_2) \gg k \& \{l_j : T_{2j}\}_{j \in J}$ and $J \subseteq I$ and $\forall j \in J. (T_{1j}, T_{2j}) \in \mathfrak{R}$.
- If $T_1 = \mu \mathbf{t}. T$, then $(\text{unfold}^1(T_1), T_2) \in \mathfrak{R}$.

where a type simulation of $(U_1, U_2) \in \mathfrak{R}$ is defined as the standard bisimulation (since U is invariant).³ The coinductive subtyping relation $T_1 \leq_c T_2$ (read: T_1 is an *asynchronous subtype* of T_2) is defined when there exists a type simulation \mathfrak{R} with $(T_1, T_2) \in \mathfrak{R}$.

An output of T_1 can be simulated after applying asynchronous optimisation \gg to the unfolded T_2 . We also need to ensure object type U_1 is a subtype of U_2 . For the input, we ensure U_2 is a subtype of U_1 . The definitions of selection and branching subsume the traditional session branching/selection subtyping.⁴ In selection a label appearing in T_1 must be included in T_2 ; dually, in branching a subtype T_1 must cover all branches declared in T_2 . For a value type, $U_1 \leq_c U_2$ implies $U_2 \leq_c U_1$ by definition. We show examples to justify our subtyping.

Example 3. Below we write $k!$ for $k!\langle U \rangle$ and $k?$ for $k?\langle U \rangle$, omitting U .

1. Let $T_1 = \mu \mathbf{t}. k?; k'!; \mathbf{t}$, $T_2 = \mu \mathbf{t}. k'!; k?; \mathbf{t}$. Then we can prove $T_2 \leq_c T_1$ using the simulation $\mathfrak{R} = \{(T_2, T_1), (k'!; k?; T_2, T_1), (k?; T_2, k?; T_1)\}$. T_2 represents more optimal communications than T_1 since it can output messages at k' without waiting.
2. Let $T'_2 = k'!; T_1$ which means first sending a signal at k' then repeating input-output actions. Then $T'_2 \leq_c T_1$ by taking $\mathfrak{R} = \{(T'_2, T_1), (T_1, k?; T_1), (k?; T'_2, k?; T_1)\}$ as a simulation closure. Note also $T_2 \leq_c T'_2$ and $T'_2 \leq_c T_2$.

² In [10], $\text{unfold}(T)$ repeatedly unfolds consecutive top-level recursion until a guarded type is obtained. In our definition, $\text{unfold}^1(T)$ expands a single recursion, not only top-level but also under guarded types.

³ Note that G is invariant like standard channel types $\tilde{[T]}$ [10].

⁴ We follow the subtyping relation in [7, 13] whose ordering is reversed from [10] since in our judgement the session environment is declared on the right-hand side of a process.

3. Let $T'_4 = k_1!; k_2!; T_3$ with $T_3 = \mu \mathbf{t}. k_3?; k_1!; k_4?; k_2!; \mathbf{t}$ and $T_4 = \mu \mathbf{t}. k_1!; k_3?; k_2!; k_4?; \mathbf{t}$. These types are extended from T'_2, T_1 and T_2 with two signal messages at the top level. Then $T'_4 \leq_c T_3$. To simulate T'_4 , we require nested unfold for T_3 . More exactly, the intermediate type $k_1!; k_4?; k_2!; T_3$ can be simulated by $k_4?; T_3$ if T_3 unfolds and $k_1!$ under recursion appears at the top-level. Similarly for $T'_4 \leq_c T_4$.
4. Take $T_5 = \mu \mathbf{t}_1. k_1!; \mu \mathbf{t}_2. k_1!; \& \{l_1 : k_2?; k_1!; \mathbf{t}_1, l_2 : k_1!; \mathbf{t}_2\}$ and let $T_6 = \mu \mathbf{t}_1. \mu \mathbf{t}_2. \& \{l_1 : k_2?; k_1!; \mathbf{t}_1, l_2 : k_1!; \mathbf{t}_2\}$. Then $T_5 \leq_c T_6$. This example is proved similarly to (3).

Note that none of the above subtyping relations, except $T_2 \leq_c T'_2$ and $T'_2 \leq_c T_2$, can be derived without including \ll in the typed simulation.

Before we prove that \leq_c is a preorder, we show that there are connecting simulations relating the components of two subtyping relations. We write $T_1 \mathfrak{R}_1 T_2$ for $(T_1, T_2) \in \mathfrak{R}_1$.

Lemma 4. *If $T_1 \mathfrak{R}_1 T_2$ and $T_2 \mathfrak{R}_2 T_3$ for type simulations \mathfrak{R}_1 and \mathfrak{R}_2 then there exists a type simulation \mathfrak{R}_3 such that if $\text{unfold}^n(T_2) \gg T'_2$, then $T'_2 \mathfrak{R}_3 T_3$.*

Definition 5 (Transitivity connection). For simulations \mathfrak{R}_1 and \mathfrak{R}_2 , we say \mathfrak{R}_3 (from the condition in Lemma 4) is a transitivity connection of $T_1 \mathfrak{R}_1 T_2$ and $T_2 \mathfrak{R}_2 T_3$. We write $\mathbf{trc}(T_1 \mathfrak{R}_1 T_2 \mathfrak{R}_2 T_3)$ for \mathfrak{R}_3 . We define $\mathbf{trc}(\mathfrak{R}_1, \mathfrak{R}_2)$ as the smallest relation such that if $(T_1, T_2) \in \mathfrak{R}_1$ and $(T_2, T_3) \in \mathfrak{R}_2$, then $\mathbf{trc}(T_1 \mathfrak{R}_1 T_2 \mathfrak{R}_2 T_3) \subseteq \mathbf{trc}(\mathfrak{R}_1, \mathfrak{R}_2)$.

From the definition, $\mathbf{trc}(\mathfrak{R}_1, \mathfrak{R}_2)$ only contains type simulations, and as the union of these it is also a type simulation. Note that the smallest relation exists, by set inclusion of relation pairs, containing all the transitivity connections of elements in $\mathfrak{R}_1/\mathfrak{R}_2$. For example, $\mathfrak{R}_3 = \mathbf{trc}(T_1 \mathfrak{R}_1 T_2 \mathfrak{R}_2 T_3)$ does not contain $(k! \langle U \rangle; \text{end}, k? \langle U \rangle; \text{end})$, which cannot be a member of any type simulation; and by set inclusion, it is smaller than $\mathfrak{R}_3 \cup (k! \langle U \rangle; \text{end}, k? \langle U \rangle; \text{end})$.

Theorem 6. *The relation \leq_c is a preorder.*

Proof. Using as standard the relation $\{(T, T) \mid T \in \text{Type}\}$, we prove \leq_c is reflexive. For transitivity of \leq_c , suppose $T_1 \leq_c T_2 \leq_c T_3$ and let \mathfrak{R}_1 and \mathfrak{R}_2 be type simulations with $(T_1, T_2) \in \mathfrak{R}_1$ and $(T_2, T_3) \in \mathfrak{R}_2$. To show $T_1 \leq_c T_3$ we need to find a type simulation \mathfrak{R} such that $(T_1, T_3) \in \mathfrak{R}$. Define \mathfrak{R} as $(\mathfrak{R}_1 \cdot \mathfrak{R}_2) \cup (\mathfrak{R}_1 \cdot \mathbf{trc}(\mathfrak{R}_1, \mathfrak{R}_2))$. Clearly $(T_1, T_3) \in \mathfrak{R}$, and it remains to show that \mathfrak{R} is a type simulation. For any $(T, T'') \in \mathfrak{R}$, there are two cases (relations above), and six subcases (simulation rules). For $(U, U') \in \mathfrak{R}$, the result is easy as U types are invariant. We only show one of the most interesting cases.

Suppose $(T, T'') \in \mathfrak{R}_1 \cdot \mathfrak{R}_2$ and $T = k! \langle U_1 \rangle; T_1$. Then there exists $(T, T') \in \mathfrak{R}_1$ and $(T', T'') \in \mathfrak{R}_2$. By the definition of type simulation, we have $\text{unfold}^n(T') \gg k! \langle U'_1 \rangle; T'_1$ and $(U_1, U'_1) \in \mathfrak{R}_1$ and $(T_1, T'_1) \in \mathfrak{R}_1$. Let $\mathbf{trc}(T \mathfrak{R}_1 T' \mathfrak{R}_2 T'') = \mathfrak{R}_3 \subseteq \mathbf{trc}(\mathfrak{R}_1, \mathfrak{R}_2)$, then by Lemma 4 we obtain $(k! \langle U'_1 \rangle; T'_1, T'') \in \mathfrak{R}_3$, and by the definition of simulation we have $\text{unfold}^m(T'') \gg k! \langle U''_1 \rangle; T''_1$ and $(U'_1, U''_1) \in \mathfrak{R}_3$ and $(T'_1, T''_1) \in \mathfrak{R}_3$. Finally, by the definition of $\mathfrak{R}_1 \cdot \mathbf{trc}(\mathfrak{R}_1, \mathfrak{R}_2)$, $(U_1, U''_1) \in \mathfrak{R}$ and $(T_1, T''_1) \in \mathfrak{R}$ as required. Other cases are similar. \square

3.3 Algorithmic Asynchronous Subtyping

The algorithmic subtyping of session types is studied in [10, § 5.1]. Due to the incorporation of asynchronous permutation and n -time unfolding in the type simulation, we

need the bound of unfolding for constructing a terminating algorithm. We first list some selected rewriting rules \xrightarrow{k} which move the types with channel k to the head applying the rules of \gg in Figure 1 in the reverse direction.

$$\begin{array}{c}
\text{(OI)} \quad k'?\langle U' \rangle; k!\langle U \rangle; T \xrightarrow{k} k!\langle U \rangle; k'?\langle U' \rangle; T \quad \text{(Tr)} \quad \frac{T_1 \xrightarrow{k} T_2 \quad T_2 \xrightarrow{k} T_3}{T_1 \xrightarrow{k} T_3} \\
\text{(CO)} \quad \frac{T \xrightarrow{k} T'}{k'!\langle U \rangle; T \xrightarrow{k} k'!\langle U \rangle; T'} \quad \text{(CB)} \quad \frac{T_j \xrightarrow{k} T'_j}{k' \& \{l_1 : T_1, \dots, l_j : T_j, \dots\} \xrightarrow{k} k' \& \{l_1 : T_1, \dots, l_j : T'_j, \dots\}}
\end{array}$$

We omit the similar rules for (OB–SS), (CI,CS), which are defined similarly to (OI) and (CO,CB). Note that we do not define \xrightarrow{k} for (E) and (M). (CO,CB) are for congruency. For a simple example, let $T_0 = k \oplus \{l_1 : k_1?\langle U_1 \rangle; k_2!\langle U_2 \rangle; \text{end}, l_2 : k_2!\langle U_2 \rangle; \text{end}\}$. Then $T_0 \xrightarrow{k_2} k \oplus \{l_1 : k_2!\langle U_2 \rangle; k_1?\langle U_1 \rangle; \text{end}, l_2 : k_2!\langle U_2 \rangle; \text{end}\} \xrightarrow{k_2} k_2!\langle U_2 \rangle; k \oplus \{k_1?\langle U_1 \rangle; \text{end}, l_2 : \text{end}\}$ by (CS,OS). We can easily show \xrightarrow{k} is confluent and terminates, and $T \xrightarrow{k} T'$ implies $T' \ll T$. We can also prove if $T \ll T'$, then we always have $T' \xrightarrow{k_1} \dots \xrightarrow{k_n} T$ where $k_1 k_2 \dots k_n$ are a (possibly empty) subsequence of channels occurring at the top-level in T with this order (e.g., $k_1 k_2 k_3 k_4$ if $k_1!; k_2 \oplus \{l_1 : k_3!, l_2 : k_4!\}$). Hence:

Proposition 7. *Given T and T' , $T \ll T'$ is decidable.*

The derivability of judgement $\Sigma \vdash T \leq T'$ is defined in Figure 2 where Σ is a sequence of assumed instances of the subtyping relation. We use n -hole type contexts $(\mathcal{T}, \mathcal{T}', \dots)$ where $[\]^{h \in H}$ denotes a hole with index h .

$$\mathcal{T} ::= [\]^{h \in H} \mid k!\langle U \rangle; \mathcal{T} \mid k'?\langle U \rangle; \mathcal{T} \mid k \oplus \{l_i : \mathcal{T}_i\}_{i \in I} \mid k \& \{l_i : \mathcal{T}_i\}_{i \in I}$$

For example, with $H = \{1, 2\}$ and $\mathcal{T} = k \oplus \{l_1 : k_1?\langle U_1 \rangle; [\]^{1 \in H}, l_2 : [\]^{2 \in H}\}$, we have $\mathcal{T}[T_i]^{i \in H} = k \oplus \{l_1 : k_1?\langle U_1 \rangle; T_1, l_2 : T_2\}$. A hole in \mathcal{T} does not appear under recursion since \ll permutes top-level actions only. We also use (1) function $\text{top}(T)$ which returns the channel at the head of T and (2) function $\text{depth}(k, T)$ to calculate how many unfoldings are needed for k to appear at the top-level. If k does not appear in T , $\text{depth}(k, T)$ is undefined. When $\text{depth}(k, T)$ is defined, $\text{depth}(k, T)$ is finite.

$$\begin{array}{l}
\text{top}(\text{end}) = \bullet \quad \text{top}(k'?\langle U \rangle; T) = \text{top}(k!\langle U \rangle; T) = \text{top}(k \& \{l_i : T_i\}_{i \in I}) = \text{top}(k \oplus \{l_i : T_i\}_{i \in I}) = k \\
\text{depth}(k, T) = 0 \quad \text{if } \text{top}(T) = k \quad \text{depth}(\bullet, \text{end}) = 0 \\
\text{depth}(k, k'?\langle U \rangle; T) = \text{depth}(k, k'!\langle U \rangle; T) = \text{depth}(k, T) \quad k \neq k' \\
\text{depth}(k, k' \& \{l_i : T_i\}_{i \in I}) = \text{depth}(k, k' \oplus \{l_i : T_i\}_{i \in I}) = \max_{i \in I}(\text{depth}(k, T_i)) \quad k \neq k' \\
\text{depth}(k, \mu t.T) = \text{depth}(k, T[\mu t.T/t]) + 1 \quad \text{depth}(\bullet, \mu t.T) = \text{depth}(\bullet, T[\mu t.T/t]) + 1
\end{array}$$

In Figure 2, [ASMP,END] are standard. In [OUT], we fix the subtype and apply \xrightarrow{k} to place $k!\langle U \rangle$ to the top level. Then we check the tail of the result of rewriting $\mathcal{T}[T_{2h}]^{h \in H}$ is a supertype of T_1 (the rule subsumes the case $k!\langle U \rangle$ already at the top level). Rule [SEL] is similarly defined. Rule [RECL] is standard. Rule [RECR] unfolds T' until a type with the same channel as the top of T appears at the top-level. The rule for input/branching is defined like [OUT]/[SEL], respectively.

$$\begin{array}{c}
\text{[ASMP]} \frac{T \leq T' \in \Sigma}{\Sigma \vdash T \leq T'} \quad \text{[END]} \frac{}{\Sigma \vdash \text{end} \leq \text{end}} \\
\text{[OUT]} \frac{\Sigma \vdash U_1 \leq U_2 \quad \Sigma \vdash T_1 \leq \mathcal{S}[T'_{2h}]^{h \in H} \quad \mathcal{S}[k!(U_2); T_{2h}]^{h \in H} \xrightarrow{k} k!(U_2); \mathcal{S}[T'_{2h}]^{h \in H}}{\Sigma \vdash k!(U_1); T_1 \leq \mathcal{S}[k!(U_2); T_{2h}]^{h \in H}} \\
\text{[SEL]} \frac{\forall i \in I. \Sigma \vdash T_i \leq \mathcal{S}[T''_{ih}]^{h \in H} \quad \mathcal{S}[k \oplus \{l_i : T'_{ih}\}_{i \in J}]^{h \in H} \xrightarrow{k} k \oplus \{l_i : \mathcal{S}[T''_{ih}]^{h \in H}\}_{i \in J} \quad I \subseteq J}{\Sigma \vdash k \oplus \{l_i : T_i\}_{i \in I} \leq \mathcal{S}[k \oplus \{l_i : T'_{ih}\}_{i \in J}]^{h \in H}} \\
\text{[RECL]} \frac{\Sigma, \mu t. T \leq T' \vdash \text{unfold}^1(\mu t. T) \leq T'}{\Sigma \vdash \mu t. T \leq T'} \quad \text{[RECR]} \frac{n = \text{depth}(\text{top}(T), T') \quad n \geq 1}{\Sigma, T \leq T' \vdash T \leq \text{unfold}^n(T')} \quad \Sigma \vdash T \leq T'
\end{array}$$

Fig. 2. Algorithmic Subtyping Rules

The rules give an algorithm for checking *the algorithmic subtyping relation* \leq (by reading these rules from upwards). As usual, [ASMP] should always be used if it is applicable, and when both [RECL] and [RECR] are applicable, [RECL] is used in preference to [RECR]. Similarly, other rules are applied in preference to [RECR], which can only be applied if the top of T does not appear at the top level of T' . As an example, let $T_1 = k \oplus \{l_1 : k_1? \langle U_1 \rangle; \text{end}\}$. Then we can derive $k_2! \langle U_2 \rangle; T_1 \leq T_0$ (T_0 is given above) by using [OUT]. At the top level, the algorithm is applied to the initial goal $\emptyset \vdash T \leq T'$ (which we often write $T \leq T'$).

Lemma 8. 1. *The subtyping algorithm always terminates.*

2. *If $T \leq_c T'$ then the algorithm does not return false when applied to $\Sigma \vdash T \leq T'$.*

The proof uses techniques related to those developed in [10]; the main differences are, for the proof of (1), we have to take the subterms up to \ll with the finite number of unfolding when we argue the size of Σ cannot increase without bound. However since \ll does not change the size of the judgement (defined in [10, Lemma 10]), we can prove (1). The proof of (2) is standard from (1).

Theorem 9 (Soundness and Completeness of the Algorithmic Subtyping). *For all closed types T and T' , $T \leq_c T'$ if and only if $T \leq T'$.*

The if-direction is by Lemma 8 (2) and the only-if direction by constructing a relation following [10, Theorem 4].

3.4 Local Asynchronous Commutative Session Typing

The type judgement for end-point processes is of the shape $\Gamma \vdash P \triangleright \Delta$ which reads: “under the environment Γ , process P has typing Δ ” where environments are defined as:

$$\Gamma ::= \emptyset \mid \Gamma, u : U \mid \Gamma, X : \tilde{U} \tilde{T} \quad \Delta ::= \emptyset \mid \Delta, \tilde{s} : \{T_p @ p\}_{p \in I}$$

A *sorting* (Γ, Γ', \dots) is a finite map from names to value types and from process variables to sequences of value types and session types. *Typing* (Δ, Δ', \dots) records linear usage of session channels. $T @ p$ is called *located type* which means T is a session type of the

participant p . In multiparty sessions, it assigns a family of located types to a vector of session channels. The typing system is identical with [13]: we only have to add the subsumption rule: i.e. $\Gamma \vdash P \triangleright \Delta$ and $\Delta \leq \Delta'$ then $\Gamma \vdash P \triangleright \Delta'$ where $\Delta \leq \Delta'$ is defined by pointwise application of \leq .


Theorem 10 (Subject Congruence and Reduction). $\Gamma \vdash P \triangleright \emptyset$ and $P \equiv Q$ imply $\Gamma \vdash Q \triangleright \emptyset$; and $\Gamma \vdash P \triangleright \emptyset$ and $P \longrightarrow Q$ imply $\Gamma \vdash Q \triangleright \emptyset$.

The proof follows the same routine as in [13], but we must take care that all permutations defined by \ll do not affect the input-output causal dependencies of the global types. We can also obtain the other three key properties, communication-safety, session-fidelity and progress as stated in [13, § 5]. The rest of the paper can be read without knowing the details of a typing system.

4 Principal Global Typing through Graph-Based Types

Why graph-based types. Let $P \stackrel{\text{def}}{=} a[p](\bar{s}).s_1!(\bar{3});s_2?(x)$ and $Q \stackrel{\text{def}}{=} \bar{a}[p](\bar{s}).s_2!(\text{true});s_1?(y)$ where P is a participant named by p and Q is an initiator named by q . Then P and Q are typable under essentially only two global types, $G = p \rightarrow q : k\langle \text{int} \rangle; q \rightarrow p : h\langle \text{bool} \rangle; \text{end}$ and $G' = q \rightarrow p : h\langle \text{bool} \rangle; p \rightarrow q : k\langle \text{int} \rangle; \text{end}$. Note the projection of G to p is \leq -minimal for P (i.e. other local types of P can be derived by subsumption): but this is not true for its projection to q . Similarly G' does not give a minimal type for P . Thus there is no “best” global type for $P|Q$: this is because interactions between P and Q take place in a criss-crossing way: the syntax of global types, which can only represent tree-like causality, is too rigid to represent such a situation.

Local and global graphs. A *local graph* is a (finite or infinite) directed graph where each node, called *action*, is labelled by one of $k?\langle U \rangle$ (input), $k!\langle U \rangle$ (output), $k\&[l_i]_{i \in I}$ (branching), $k \oplus [l_i]_{i \in I}$ (selection) and $k \oplus l$ (label-output [1]); and, for edges: (1) each edge from $k\&[l_i]_{i \in I}$ or $k \oplus [l_i]_{i \in I}$ is labelled by one of $\{l_i\}$; and (2) $k \oplus l$ (resp. $k!\langle U \rangle$) has a unique outgoing edge, and its target is always an output/selection/label-output at k . A *global graph for participants* $\{p_1, \dots, p_n\}$, written $\mathcal{G}, \mathcal{G}', \dots$, is a disjoint union of an $\{p_1, \dots, p_n\}$ -indexed family of local graphs. Given \mathcal{G} for $\{p_1, \dots, p_n\}$, its p_i -*component* is the local graph in \mathcal{G} indexed by p_i . A node is *active* if it has no incoming edges.



In (a) above, we show a global graph for P and Q given above, consisting of two local graphs (balloons labelled by p and q), each with an output, an input and no edges. If we add an edge from input to output in each local graph, we get the global graph (b) for $a[p](\bar{s}).s_2?(x);s_1!(\bar{3})$ and $\bar{a}[p](\bar{s}).s_1?(y);s_2!(\text{true})$, which now deadlocks.

Linearity, progress and coherence. We equip global graphs with a notion of reduction which abstracts that of processes. Below we write $\mathcal{E}[\cdot][\cdot][\cdot]$ for a global graph with one or more holes, each of which is to be filled with a sub-graph of a local graph, such that all holes are active, i.e. have no incoming edges.

$$\begin{aligned} \mathcal{E}[k?\langle U \rangle][k!\langle U \rangle] &\longrightarrow \mathcal{E}[\emptyset] & \mathcal{E}[k\&[l_i : \mathcal{G}_i]_{i \in I}][k \oplus l_j] &\longrightarrow \mathcal{E}[\mathcal{G}_j][\emptyset] \quad (j \in I) \\ & & \mathcal{E}[k \oplus [l_i : \mathcal{G}_i]_{i \in I}] &\longrightarrow \mathcal{E}[k \oplus l_i; \mathcal{G}_i] \end{aligned}$$

Above \emptyset is the empty graph. In each rule, the replacement in the hole(s) entails taking off *both* the old graph *and* all the outgoing edge(s) from it and filling the hole with the new graph. A reduction by the first two rules is called *communication at k*. In the second rule, $k\&[l_i : \mathcal{G}_i]_{i \in I}$ is the disjoint union of $k\&[l_i]_{i \in I}$ and $\{\mathcal{G}_i\}_{i \in I}$ together with, for each $i \in I$, l_i -labelled edges from $k\&[l_i]_{i \in I}$ to all the active actions in \mathcal{G}_i . In the third (from [1]), $k \oplus [l_i : \mathcal{G}_i]_{i \in I}$ is as $k\&[l_i : \mathcal{G}_i]_{i \in I}$ while $k \oplus l_i; \mathcal{G}_i$ is the disjoint union of $k \oplus l_i$ and \mathcal{G}_i with edges from the former to the active output/selection/label-outputs at k in \mathcal{G}_i . A global graph \mathcal{G} is *linear* when for each \mathcal{G}' such that $\mathcal{G} \longrightarrow^* \mathcal{G}'$, if \mathcal{G}' has two active actions at k , a reduction at k is possible, and no other active action shares k . A global graph \mathcal{G} has *progress* when for each \mathcal{G}' such that $\mathcal{G} \longrightarrow^* \mathcal{G}'$, either \mathcal{G}' reduces or it is empty. Finally we say \mathcal{G} is *coherent* when it is linear and has progress.

Coherent global graphs from local types. A local graph is constructed from a local type as the latter's regular tree representations. Given $\Delta = \{T_i @ p_i\}_{i \in I}$, this immediately gives the global graph for $\{p_i\}_{i \in I}$, which we write $\llbracket \Delta \rrbracket$. The coherence of $\llbracket \Delta \rrbracket$ is decidable, as we outline below.

We first check Δ is *well-directed* in the sense that each channel in Δ is used by two and only two participants and moreover one of them uses it only for input/branching and the other only for output/selection/label-output, which can be checked by going through Δ once. For well-directed Δ , there is an algorithm to ensure linearity of $\llbracket \Delta \rrbracket$, by checking if each pair of participants in Δ are compatible in their type structures, closely following the algorithmic subtyping in § 3.2.

Through the validation of compatibility of Δ , we can equip $\llbracket \Delta \rrbracket$ with the additional *communication edges*, from each output/selection/label-output to its potentially interacting action(s), representing potential redexes. Using this added set of edges, we reduce the progress of $\llbracket \Delta \rrbracket$ to the acyclicity of its paths consisting of its local and communication edges, completely characterising progress under linearity. The acyclicity of $\llbracket \Delta \rrbracket$ is then reducible to that of its initial finite sub-graph. Because linearity and compatibility are equivalent under progress, we obtain:

Theorem 11 (complete algorithmic characterisation of coherence). *Let $\Delta = \{T_i @ p_i\}_{i \in I}$ be well-directed. Then the coherence of $\llbracket \Delta \rrbracket$ with Δ given as input is decidable.*

Principal global typing through global graphs. Any projectable global type G for participants say $\{p_i\}_{i \in I}$ is equivalent to its projections $\Delta = \{(G \upharpoonright p_i) @ p_i\}_{i \in I}$, and because such Δ is immediately compatible and acyclic, we can regard G as a coherent global graph. This motivates the use of coherent global graphs instead of global types

in the type discipline, presenting $\llbracket \Delta \rrbracket$ as Δ itself.⁵ By replacing global types with coherent global graphs in types and typing rules, we obtain a new type discipline. We write $\Gamma \vdash_g P \triangleright \Delta$ for typability in this new discipline (subsumption is consistent because if $\llbracket \Delta \rrbracket$ is coherent and Δ' is point-wise \leq -smaller than Δ then $\llbracket \Delta' \rrbracket$ is also coherent). By identifying \mathcal{G} as the corresponding coherent global graph, $\Gamma \vdash P \triangleright \Delta$ implies $\Gamma \vdash_g P \triangleright \Delta$. Further, since linearity and progress of $\llbracket \Delta \rrbracket$ are reflected onto the dynamics of typed processes (precisely following the arguments in [13]), the typability in \vdash_g ensures communication safety and progress.

For the principal typing property, we add the \leq -least element \perp to the set of local types; \perp is also used as local graph occurring in global graphs (where intuitively \perp denotes a placeholder for a local behaviour). The coherence and other notions for global graphs are defined ignoring \perp . Without loss of practical generality we assume each shared name say a has a fixed *arity* which is the number of participants for a potential session established through a ; and that processes are type-annotated on bound variables and free object names in the standard way. Through local type inference [6, 15] using the point-wise join of coherent global types (calculated as in algorithmic subtyping), together with Theorem 11, we obtain a principal global typing property. Below we write $\Gamma' \leq \Gamma$ for $\text{dom}(\Gamma') \subset \text{dom}(\Gamma)$ and $\Gamma'(a) \leq \Gamma(a)$ for each $a \in \text{dom}(\Gamma')$. We say P is *closed* if it has no free session channels nor free variables.

Theorem 12 (principal global typing). *Let P be closed. (1) The typability of P with respect to \vdash_g is decidable. (2) If P is typable then P has a principal global typing Γ_0 in the sense that $\Gamma_0 \vdash_g P \triangleright \emptyset$ holds and moreover $\Gamma \vdash_g P \triangleright \emptyset$ implies $\Gamma_0 \leq \Gamma$.*

5 Application: Double-Buffering Algorithm

This section illustrates the use of our type theory using the *double-buffering algorithm* [21], a basic distributed algorithm widely used in stream/media processing and high-performance and multicore computing, presenting how the two strategies discussed in Introduction can be applied through the theories presented in the previous sections.

The purpose of the double-buffering algorithm is to transform a large amount of data, where a series of chunks of data are transferred from a source (Source) to a transformer (called Kernel), gets processed there and delivered to a sink (Sink). Under potential temporal variations in processing and communication time, it is necessary to synchronise among these three parties through message passing. However a naive, and obviously safe, protocol leads to a highly sequential, non-optimal distributed algorithm. Thus it is beneficial to increase asynchrony of local programs without violating the shared protocol. We show the outline of an application of our theories to achieve this goal, starting from a sequential and safe global protocol to optimised local protocols through asynchronous communication subtyping, with a formal safety guarantee.

⁵ To be precise, we regard Δ up to the type isomorphism corresponding to \leq ; and we take off, from each branching type, its branches (if any) which never get invoked in any reduction path: such “garbage” branches are precisely identified during the validation of coherence.

<p><i>Global Type</i> : $G =$</p> $\begin{aligned} &\mu\mathbf{t}.(\\ &\text{K} \rightarrow \text{So} : r_1 \langle \rangle; \quad \text{K} \rightarrow \text{So} : r_2 \langle \rangle; \\ &\text{So} \rightarrow \text{K} : s_1 \langle U \rangle; \quad \text{So} \rightarrow \text{K} : s_2 \langle U \rangle; \\ &\text{Si} \rightarrow \text{K} : t_1 \langle \rangle; \quad \text{Si} \rightarrow \text{K} : t_2 \langle \rangle; \\ &\text{K} \rightarrow \text{Si} : u_1 \langle U \rangle; \quad \text{K} \rightarrow \text{Si} : u_2 \langle U \rangle; \mathbf{t} \end{aligned}$ <p><i>Projected Local Type of Kernel</i> :</p> $\begin{aligned} T = \\ &\mu\mathbf{t}.r_1 \langle \rangle; s_1 \langle U \rangle; t_1 \langle \rangle; u_1 \langle U \rangle; \\ &\quad r_2 \langle \rangle; s_2 \langle U \rangle; t_2 \langle \rangle; u_2 \langle U \rangle; \mathbf{t} \end{aligned}$ <p><i>Local Type of Kernel</i> :</p> $\begin{aligned} T^* = \\ &r_1 \langle \rangle; r_2 \langle \rangle; \\ &\mu\mathbf{t}.s_1 \langle U \rangle; t_1 \langle \rangle; u_1 \langle U \rangle; r_1 \langle \rangle; \\ &\quad s_2 \langle U \rangle; t_2 \langle \rangle; u_2 \langle U \rangle; r_2 \langle \rangle; \mathbf{t} \end{aligned}$	<p><i>Source</i> :</p> $\begin{aligned} &a[1](r_1 r_2 s_1 s_2 t_1 t_2 u_1 u_2). \\ &\mu X.(\\ &\quad \dots // \text{assign data to } y[1..n] \\ &\quad r_1 \langle \rangle; s_1 \langle y[1..n] \rangle; \\ &\quad \dots // \text{assign data to } y[1..n] \\ &\quad r_2 \langle \rangle; s_2 \langle y \rangle; X) \end{aligned}$ <p><i>Sink</i> :</p> $\begin{aligned} &a[2](r_1 r_2 s_1 s_2 t_1 t_2 u_1 u_2). \\ &\mu X.(\\ &\quad t_1 \langle \rangle; u_1 \langle z \rangle; \\ &\quad \dots // \text{print } z[1..n] \\ &\quad t_2 \langle \rangle; u_2 \langle z \rangle; \\ &\quad \dots // \text{print } z[1..n] \\ &\quad X) \end{aligned}$	<p><i>Kernel</i> :</p> $\begin{aligned} &\bar{a}[1,2](r_1 r_2 s_1 s_2 t_1 t_2 u_1 u_2). \\ &r_1 \langle \rangle; r_2 \langle \rangle; \\ &\mu X.(\\ &\quad s_1 \langle x_A \rangle; \\ &\quad \dots // \text{repeat:} \\ &\quad \dots // x_A[i] ::= x_A[i] \oplus \text{key} \\ &\quad \dots // \text{key} ::= x_A[i] \\ &\quad t_1 \langle \rangle; u_1 \langle x_A \rangle; r_1 \langle \rangle; \\ &\quad s_2 \langle x_B \rangle; \\ &\quad \dots // \text{repeat:} \\ &\quad \dots // x_B[i] = x_B[i] \oplus \text{key} \\ &\quad \dots // \text{key} = x_B[i] \\ &\quad t_2 \langle \rangle; u_2 \langle x_B \rangle; r_2 \langle \rangle; X) \end{aligned}$
--	---	---

Fig. 3. Double-Buffering Algorithm: Processes and Types

Top-down approach (1): global type. The development of programs starts from the global type G on the left-most column in Figure 3. So, K and Si denote participant names for Source, Kernel and Sink. U denotes a large int-array type. Assuming Kernel will use two channels and the associated arrays for potential parallelism, the global type G starts from a recursion, describing an infinite loop. In the loop, Kernel first notifies Source via $r_{1,2}$ that it is ready to receive data in its two channels ($s_{1,2}$, with signal at r_i saying s_i is ready); Source complies, sending two chunks of data sequentially via $s_{1,2}$. Then Kernel (internally processes data and) waits for Sink to inform (via $t_{1,2}$) that Sink is ready to receive data via $u_{1,2}$: upon receiving the signals, Kernel sends the two chunks of processed data to Sink. This protocol is sequential but is safe and deadlock-free.

Top-down approach (2): local type and its refinement. Just below the global type G , Figure 3 gives the local type T of Kernel as directly projected from the global type. Our purpose is to refine T so that (1) the new local protocol is more asynchronous, allowing overlap of communication and computation [9, 11]; and (2) it still conforms to G — Kernel with the new optimised protocol will safely interact with Source and Sink who conform to the original global type G . For this purpose the developer may come up with a more asynchronous T^* , given in Figure 3 after T . In this refined protocol, Kernel notifies Source via both $r_{1,2}$, but only once before entering the loop, allowing Source to start its work. Now inside the loop, the refined protocol dictates Kernel first receives data via its first channel s_1 with Source, processes the data and sends out the result to Sink via its first channel u_1 with Sink and *immediately notifies Source via r_1 that it's ready in its first channel*, allowing Source to start sending data early. Kernel then repeats the same work for its second channels with Source and Sink. In this way, Kernel can process data it has already received in one channel while it is receiving data in the other, noting it can take time for large data to sent, transferred and received.

We now show this optimised local protocol is safe w.r.t. other participants conforming to G , through the asynchronous communication subtyping. The justification uses

nested unfolding. We start from unfolding T once to match r_1, r_2 of T^* as $\text{unfold}^1(T) = r_1! \langle \rangle; s_1? \langle U \rangle; t_1? \langle \rangle; u_1! \langle U \rangle; r_2! \langle \rangle; s_2? \langle U \rangle; t_2? \langle \rangle; u_2! \langle U \rangle; T$. Then $r_1! \langle \rangle$ matches T^* . To simulate $r_2! \langle \rangle$ of T^* , $r_2! \langle \rangle$ is permuted by \ll . Let $T^* = r_1! \langle \rangle; r_2! \langle \rangle; T_R^*$. Thus $\text{unfold}^1(T_R^*)$ must be simulated by $T' = s_1? \langle U \rangle; t_1? \langle \rangle; u_1! \langle U \rangle; s_2? \langle U \rangle; t_2? \langle \rangle; u_2! \langle U \rangle; T$. However to simulate $r_1! \langle \rangle$ in $\text{unfold}^1(T_R^*)$, T must be unfolded again since the types in the guarded position of T' do not include $r_1! \langle \rangle$. By [RECR], it now suffices to solve the following:

$$r_1! \langle \rangle; s_2? \langle U \rangle; t_2? \langle \rangle; u_2! \langle U \rangle; r_2! \langle \rangle; T_R^* \leq s_2? \langle U \rangle; t_2? \langle \rangle; u_2! \langle U \rangle; \text{unfold}^1(T).$$

For this we apply [IN,OUT] of \leq with \ll , reaching the assumption in Σ in [ASMP].

Top-down approach (3): code development. Figure 3 depicts the skeleton of the three (final) programs which conform to the global type. All participants initiate the session at a in the first line. We only illustrate the behaviour of Kernel, considering a simple transformation for stream encryption. Kernel, after initialising its variables including the initial key value, signals to Source that its buffers are both empty, via r_1 and r_2 : then enters the main loop, where it does the following: it first receives the datum at x_A via s_1 , goes through the buffer taking the XOR element-wise with key, after which it waits for Sink's cue via t_1 (which may have already arrived asynchronously), and finally sends out the buffer content to Sink via u_1 , and tells Source it's ready at A via r_1 : then similarly works for the second buffer (given in the next column). Unlike Source and Sink, the behaviour of Kernel does not conform to the projection of G to Kernel (T): however T^* does type-abstract its behaviour directly and because $T^* \leq T$ by the argument above, Kernel does type-check under T with subsumption, hence under G .

Integration with bottom-up approach. The development process described above can be effectively and seamlessly integrated with the “bottom-up” strategy discussed in Introduction through the type inference (synthesis) of global types in § 4, which allows developers to directly refine programs and to synthesise a new global protocol reflecting the refinement, incrementally validating compatibility. This added flexibility is useful since, in actual development, programmers may often directly work on programs rather than starting from refinement of local protocols.

Further examples showing the applicability of permutation of branching/selections to parallel algorithms [14] can be found in the full version.

6 Related Work

Branching/selection subtyping in session types is first studied in [10] for binary session types. We use their syntactic approach for defining a type-simulation, but a significant extension from their technique is needed due to the incorporation of \ll and nested unfoldings, which makes the proof of transitivity delicate and challenging. An initial idea of asynchronous communication subtyping for binary sessions is presented in an unpublished manuscript [18], where the treatment for recursive types and branching/selection

types is left open. A recent work in a technical report [17] demonstrates a subtyping rule similar to our (OI) rule is useful for an object calculus with asynchronous binary sessions, with an iso-recursive system. It is an interesting future work to extend to the $HO\pi$ -calculus [16] where a careful formulation for the algorithmic subtyping would be required in the presence of arrow types. The top-down approach in multiparty session types is first studied in [13], but a local refinement (asynchronous subtyping) is not proposed there. The problem of synthesising a global specification from endpoint behaviours has been a lingering question since the inception of the notion of global descriptions for business protocols [24], being posed as an open problem in [1, 2, 13]. Inference of principal types is studied in [15] for binary session types (note in binary sessions the issue of global synthesis does not arise). The present work gives clear and general solutions to these extant technical problems.

In the context of multiparty session types, a typing system for a strong progress property is studied in [1]. Asynchronous communication subtyping can be smoothly applied to [1]. For delegations, the main definitions of \ll , \leq_c and \leq stay as the same, but proofs need to be revised to treat nesting types; and for the principal typing, \perp should be added into a carried type in global graphs. The study of formal theories of contracts are studied in [8] using CCS-like processes as a type representation. The work [19] extends [8] with the treatment of asynchronous behaviours using orchestrators, through the use of bounded buffers that control message flows between a client and servers. Our own system in [7] developed a theory in which a global specification arises as a programming language itself.

Conformance and refinement based on agreement of service specifications is studied in [3], using a synchronous CCS-based calculus as a contract language, and testing-preorders to check subcontract compliance. Neither type-checking of end-point processes using projected contracts (in our case, Theorem 9) nor a bottom-up strategy is presented there. The work [5] proposes a distributed calculus with sessions, incorporating the merging of running sessions. Another work [23] presents a calculus for service orientations by extending the π -calculus with context-sensitive interactions, equipped with service and request primitives and local exceptions. These preceding works do not treat the main technical problems addressed in the present work – the asynchronous communication subtyping, type-based local refinement/conformance, and a derivation of the minimum global types, backed-up by the efficient type-checking and inference algorithms, ensuring strong safety properties based on the session type discipline.

Acknowledgements. We thank the reviewers for their useful comments, Gary Brown and Steve Ross-Talbot for discussions on the potential applications of the presented framework for software development, and Matthew Rawlings for discussions on the practical significance of asynchrony in financial protocols. We generalised \ll with input commutativity following a suggestion by Raymond Hu. The work is partially supported by EPSRC GR/T03208, GR/T03215, EP/F002114, EP/F003757 and IST2005-015905 MOBIUS.

References

1. Bettini, L., et al.: Global progress in dynamically interleaved multiparty sessions. In: van Breugel, F., Chechik, M. (eds.) CONCUR 2008. LNCS, vol. 5201, pp. 418–433. Springer, Heidelberg (2008)
2. Bonelli, E., Compagnoni, A.: Multipoint Session Types for a Distributed Calculus. In: Barthe, G., Fournet, C. (eds.) TGC 2007. LNCS, vol. 4912, pp. 240–256. Springer, Heidelberg (2008)
3. Bravetti, M., Zavattaro, G.: A theory for strong service compliance. In: Murphy, A.L., Vitek, J. (eds.) COORDINATION 2007. LNCS, vol. 4467, pp. 96–112. Springer, Heidelberg (2007)
4. Bravetti, M., Zavattaro, G.: Towards a unifying theory for choreography conformance and contract compliance. In: Lumpe, M., Vanderperren, W. (eds.) SC 2007. LNCS, vol. 4829, pp. 34–50. Springer, Heidelberg (2007)
5. Bruni, R., Lanese, I., Melgratti, H., Tuosto, E.: Multiparty Sessions in SOC. In: Lea, D., Zavattaro, G. (eds.) COORDINATION 2008. LNCS, vol. 5052, pp. 67–82. Springer, Heidelberg (2008)
6. Carbone, M., Honda, K., Yoshida, N.: A theoretical basis of communication-centered concurrent programming. To appear as a WS-CDL working report, www.dcs.qmul.ac.uk/~carbonem/cdlpaper/workingnote.pdf
7. Carbone, M., Honda, K., Yoshida, N.: Structured communication-centred programming for web services. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 2–17. Springer, Heidelberg (2007)
8. Castagna, G., Gesbert, N., Padovani, L.: A theory of contracts for web services. In: POPL, pp. 261–272 (2008)
9. Culler, D., et al.: Logp: towards a realistic model of parallel computation. SIGPLAN Not. 28(7), 1–12 (1993)
10. Gay, S., Hole, M.: Subtyping for Session Types in the Pi-Calculus. Acta Informatica 42(2/3), 191–225 (2005)
11. Gschwind, M.: The cell broadband engine: Exploiting multiple levels of parallelism in a chip multiprocessor. International Journal of Parallel Programming 35(3), 233–262 (2007)
12. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type disciplines for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998)
13. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. In: POPL 2008, pp. 273–284. ACM, New York (2008)
14. Mattson, T., Sanders, B., Massingill, B.: Patterns for Parallel Programming. Addison Wesley, Reading (2005)
15. Mezzina, L.G.: How to infer finite session types in a calculus of services and sessions. In: Lea, D., Zavattaro, G. (eds.) COORDINATION 2008. LNCS, vol. 5052, pp. 216–231. Springer, Heidelberg (2008)
16. Mostrous, D., Yoshida, N.: Two Sessions Typing Systems for Higher-Order Mobile Processes. In: Della Rocca, S.R. (ed.) TLCA 2007. LNCS, vol. 4583, pp. 321–335. Springer, Heidelberg (2007)
17. Mostrous, D., Yoshida, N.: A Session Object Calculus for Structured Communication-Based Programming. Technical report, Imperial College London (2008), www.doc.ic.ac.uk/~mostrous
18. Neubauer, M., Thiemann, P.: Session Types for Asynchronous Communication. Universität Freiburg (2004)
19. Padovani, L.: Contract-directed synthesis of simple orchestrators. In: van Breugel, F., Chechik, M. (eds.) CONCUR 2008. LNCS, vol. 5201, pp. 131–146. Springer, Heidelberg (2008)

20. Pierce, B., Sangiorgi, D.: Typing and subtyping for mobile processes. *Journal of Mathematical Structures in Computer Science* 6(5), 409–454 (1996)
21. Sancho, J.C., Kerbyson, D.J.: Analysis of Double Buffering on two Different Multicore Architectures: Quad-core Opteron and the Cell-BE. In: *International Parallel and Distributed Processing Symposium (IPDPS)*, April 14–18. IEEE, Los Alamitos (2008)
22. Takeuchi, K., Honda, K., Kubo, M.: An interaction-based language and its typing system. In: Halatsis, C., Philokyprou, G., Maritsas, D., Theodoridis, S. (eds.) *PARLE 1994*. LNCS, vol. 817, pp. 398–413. Springer, Heidelberg (1994)
23. Vieira, H.T., Caires, L., Seco, J.C.: The conversation calculus: A model of service-oriented computation. In: Drossopoulou, S. (ed.) *ESOP 2008*. LNCS, vol. 4960, pp. 269–283. Springer, Heidelberg (2008)
24. Web Services Choreography Working Group. *Web Services Choreography Description Language*, <http://www.w3.org/2002/ws/chor/>