

Explicit Connection Actions in Multiparty Session Types

Raymond Hu^(✉) and Nobuko Yoshida

Imperial College London, London, UK
rhu@doc.ic.ac.uk

Abstract. This work extends asynchronous multiparty session types (MPST) with *explicit connection actions* to support protocols with optional and dynamic participants. The actions by which endpoints are connected and disconnected are a key element of real-world protocols that is not treated in existing MPST works. In addition, the use cases motivating explicit connections often require a more relaxed form of multiparty choice: these extensions do not satisfy the conservative restrictions used to ensure safety in standard syntactic MPST. Instead, we develop a modelling-based approach to validate MPST safety and progress for these enriched protocols. We present a toolchain implementation, for distributed programming based on our extended MPST in Java, and a core formalism, demonstrating the soundness of our approach. We discuss key implementation issues related to the proposed extensions: a practical treatment of choice subtyping for MPST progress, and multiparty correlation of dynamic binary connections.

1 Introduction

Multiparty session types (MPST) is a type systems theory for verifying message passing concurrent processes, originally developed in the π -calculus [21]. A standard top-down presentation of syntactic MPST systems consists of three layers: (1) a global specification of an asynchronous message passing protocol as a *global type*, with the participants abstracted as *roles*; (2) a syntactic *projection* to a localised view of the protocol for each role as a *local type*; which are in turn used to (3) type check the endpoint *processes* implementing the roles. A well-typed system of session endpoint implementations is guaranteed free from communication safety errors, such as unexpected message receptions and deadlocks.

In our view, the central design point of practical languages and tools based on session types is: (a) to identify a class of protocols, through the constraints of the type syntax and accompanying well-formedness conditions; such that MPST safety is indeed guaranteed by (b) (independent) verification of endpoint programs against their local projections. Much research, both multiparty and the special case of *binary* sessions, has focused on addressing (b) in various ways: extending existing languages to support static session typing (e.g., Links [32]) via pre-processing tools (Java [25, 45]), embedding into existing languages via encodings (Haskell [26, 40], Rust [27]), dynamic session typing by runtime monitoring (Python [15], Erlang [19]), hybrid (part static, part dynamic) approaches (Java [24], Scala [42], ML [36]), and code generation (MPI/C [35]).

Regarding (a), the multiparty works of the above mostly follow the core theoretical systems [12,22], where protocol well-formedness is directly derived from syntactic restrictions in conjunction with various simplifying assumptions. Unfortunately, these restrictions are too conservative for many useful patterns found in practice. For example, consider a basic, generic pattern starting with interactions between two session participants that, in *some* cases, leads to the later involvement of a third party. By contrast, the standard MPST notion of session initiation is assumed to be a single, atomic synchronisation between *all* parties (as in all of the above works), which inherently rules out any instance of this pattern. Standard MPST basically do not support protocols with *dynamic joining/leaving* of participants during a session, nor *optional* participation.

This paper. We develop an MPST toolchain to address limitations w.r.t. to (a) as discussed above, that can be readily integrated with some of the existing approaches for (b). There are two main contributions.

One is to extend MPST to support *explicit connection actions* in protocol specifications, in a manner that is closely guided by the practical motivations. Rather than a globally interconnected structure between a fixed number of participants, we consider a multiparty session as a dynamically evolving configuration of *binary* bidirectional connections that are established and closed (and possibly re-established) as the session progresses. Concretely, we extend an existing MPST-based protocol description language, Scribble [44,47]. The following is an instance of the pattern from above in our extended Scribble:

```
explicit global protocol OptionalDynamicThirdParty(role A, role B, role C) {
  hello() connect A to B; // A connects to B; sends a message labelled hello
  goodday() from B to A; // B replies to A on the established connection
  choice at A { opt1() from A to B; // A has two choices: send opt1 or opt2 to B
               greetings() connect B to C; } // B connects to C; sends greetings
  or { opt2() from A to B; } // Session ends without involving C
```

(The syntax is explained more in Sect.2.) Explicit connection actions allow MPST to better fit real-world use cases from domains such as Internet applications and Web services, where multiparty systems are often implemented over binary transports like TCP and HTTP. As we shall see in examples, many patterns involving explicit connections also require a more relaxed form of choice than in standard MPST, with mixed action kinds and destination roles.

The second aspect relates to global type *validation* in our extended MPST. The proposed extensions do not satisfy the conservative restrictions used to ensure safety in standard syntactic MPST: they allow writing additional use cases, but also introduce the potential for errors that were previously precluded.

<pre>/* Standard MPST: all roles interconnected * on session init. (Scribble default) */ global protocol P1(role A, role B, role C) { choice at A { 1() from A to B; } or { 2() from A to C; } do P1(A, B, C); }</pre>	<pre>// Explicit connection actions explicit global protocol P2(role A, role B, role C) { choice at A { 1() connect A to B; disconnect A and B; } or { 2() connect A to C; disconnect A and C; } do P2(A, B, C); }</pre>
---	--

The minimal examples above illustrate some of the issues at hand. P1 features a choice involving only A and B in one case, and A and C in the other (which is not permitted in [12, 17, 18, 22]), that is repeated continuously by the recursion (not permitted in [18]). However, P1 *does* satisfy the intuitive notion of MPST *safety* (e.g., no reception errors or deadlocks); and under an assumption of *output choice fairness*, i.e., provided A does not starve B or C of messages, P1 also satisfies MPST *progress* (otherwise, if, e.g., A talks only to B, then C remains in the session but never progresses). Using explicit connection actions, this pattern can be rewritten in P2 to satisfy both safety and progress *without* such an assumption.

Our approach is to develop a modelling-based validation for MPST protocols. Specifically, we derive a model of a global type from the *1-bounded* execution of the induced multiparty session, i.e., where the capacity of each dynamically established, asynchronous channel is limited to one message; and explicitly check the model is free of the traditional MPST safety and progress errors, as well as the additional kinds of errors introduced by our extensions, such as unexpected or duplicate (dis)connections. The key to this approach is that the characteristics of syntactic MPST can be leveraged to serve the soundness of the bounded validation; as opposed to solely relying on syntactic restrictions for outright safety. We treat output choice fairness by a structural transformation in the model construction, that reflects the underlying issue of session subtyping [37]; e.g., our validation accepts P1 (above) only if fairness is assumed.

Techniques based on “minimal asynchrony” have been employed for various purposes in related theoretical works (Sect. 5); e.g., to show the decidability of choreography realisability [4], classifying session types in the context of communicating FSMs [18], and the study of properties of half-duplex binary systems [11]. The advance of this work is to formulate the 1-bounded validation for our extended MPST; and its application in a practical toolchain, from the validation of our extended Scribble specifications to safe implementations of distributed Java endpoints. We believe that such an approach may offer a practical, uniform validation methodology for MPST-based protocols, towards incorporating further MPST extensions (e.g., [5, 6, 15, 29, 46]) together in an integrated toolchain.

2 Use Case and Overview

2.1 Use Case: Travel Agency Web Service (Revisited)

Travel Agency is one of the widely-used examples in session types literature, based on a W3C Web services choreography use case;¹ we follow the version in [1]. The basic scenario starts by a *Client* (C) initiating a session with the *Travel Agent* (A) to negotiate a product quote. The client may eventually choose to reject all quotes, ending the session; or to accept one, leading to a payment transaction between the client and a third-party *Service* (S). Although this is a natural multiparty use case, it is not actually fully supported by standard MPST. To see the potential problems, consider the following fragment from the latter part of the protocol:

¹ <https://www.w3.org/TR/2004/WD-ws-chor-reqs-20040311/> Sect. 3.1.1.

```

1  explicit global protocol TravelAgency | 17 // So far, only C and A are connected
2      (role C, role A, role S) {         | 18 aux global protocol Pay
3      connect C to A;                   | 19     (role C, role A, role S) {
4      do Nego(C, A, S);                 | 20     choice at C {
5  }                                     | 21     // C connects to S, sends pay info
6  // aux subprotocols                  | 22     pay(Str) connect C to S;
7  aux global protocol Nego             | 23     // S returns a payment reference
8      (role C, role A, role S) {       | 24     confirm(Int) from S to C;
9      choice at C {                    | 25     // C forwards the payref to A
10     query(Str) from C to A;          | 26     accept(Int) from C to A;
11     quote(Int) from A to C;         | 27   } or {
12     do Nego(C, A, S);               | 28     reject() from C to A;
13   } or {                             | 29   }
14     do Pay(C, A, S);                | 30 } // End of protocol
15 } }                                  | 31
    
```

Fig. 1. The Travel Agency choreography use case¹ using explicit connection actions.

```

choice at C { pay(Str) from C to S; confirm(Int) S to C; accept(Int) from C to A; }
              or { reject () from C to A; } // S not involved [i]
    
```

In standard MPST, the execution model is that all three roles are synchronised on session initiation, and there are no further implicit messages (e.g., no session termination handshake). Under these assumptions, the above fragment is unsafe because, in the second case, there is no way for an implementation of **S** to *locally* determine that the session is finished. Consequently, specifications in existing MPST use workarounds that are less rigorous (e.g., decomposing the protocol into separate global types, losing some of the message causalities) or less realistic/efficient (e.g., by introducing extra messages, or *delegation* [12]).

The above fragment is also not permitted as a standard MPST choice due to the *directed choice* restriction: the messages from a branch point must be sent to the *same* role in all cases (e.g., r' in the type grammar $r \rightarrow r' : \{l_i.G_i\}_{i \in I}$ [12, 22]; similarly in automata-based works [17, 18]) as a conservative element towards ensuring safety. The superficial quick fix by simply moving the `accept` message to the start of the first case is not possible in this example, because the `Int` payload of this message is intended to be the value (the payment reference `Int`) received by **C** in the preceding `confirm` message.

Explicit connection actions allow this use case to be safely captured as a single global type, as given by `TravelAgency` and its two subprotocols (`aux`) in our extended Scribble in Fig. 1. Line 1 declares the root protocol with the three roles **C**, **A** and **S**. The new `explicit` modifier means that every inter-role connection used for message passing must first be established by explicitly specified connection actions. A session starts by **C** `connect to A` (line 3), creating a bidirectional channel (e.g., TCP) between client **C** and server **A**.

We then enter the `Nego` subprotocol by the `do`-statement, with the `do` argument roles playing the target parameter roles (given the same names in this example). The `choice at C` on line 9 means **C** makes an *internal* choice between the two cases (the `or`-separated blocks), to be explicitly communicated as an

external choice to other roles as appropriate. In the first case, a message of *signature* `query(Str)` (a message with header/label `query`, and one payload value of type `Str`) is sent from `C` to `A`. `A` replies with a `quote(Int)`, and the choice is repeated by the recursive `do` on line 12. `A` and `C` thus perform the `query/quote` exchange some number of times (possibly zero, in this simplified version). Finally, in `Pay`, `C` has two further options. `C` may `connect` to `S`, thereby *dynamically* bringing `S` into the session: `C` exchanges payment details `pay(Str)` for a payment reference `confirm(Int)` with `S`, and forwards the reference to `A`. Otherwise, `C` sends a `reject` to `A`, and the session ends without involving `S`. Note that these syntactically nested choices actually amount to a single choice at `C`, between *mixed* kinds of actions to *different* roles: the `connect` to `S`, and the `sends` to `A`.

Extending MPST with explicit connection actions allows such protocols because, e.g., the `connect from C to S`, serves to delimit the scope of `S`'s involvement to the relevant choice case only. From `S`'s view, the *whole* session starts and ends, by interactions with `C`, in this one case, if the session indeed proceeds this way at run-time: while `S` remains unconnected, we can consider it as “inactive” with regards to session safety and progress. At the same time, this solution reduces the gap between MPST-based descriptions and real protocols, like Internet application RFCs, by recognising that the client/server connection actions are as important to a rigorous specification as the message passing (e.g., the STARTTLS “re-connection” in SMTP [28], and FTP’s active/passive modes [39]).

The communication model promoted by our extended MPST is *at most* one (as opposed to *exactly* one) connection between any pair of roles. Consider the following `explicit` protocol with roles `A`, `B` and `C`:

```
connect A to B; rec X { [ii]
choice at A { 1() from A to B; 2() connect B to C; disconnect B and C; continue X;
              } or { 3() from A to B; } }
```

The `disconnect` is necessary, inside the *recursion* `rec X { ...continue X; }`, to ensure there is never more than one connection between `B` and `C` (similarly in `P2` in Sect. 1). We can assume implicit `disconnect` actions at the end of a protocol.

2.2 Overview of 1-Bounded Global Type Validation and Examples

The restrictions employed in standard MPST are convenient for reasoning about the MPST safety properties. Aside from surface syntax details, systems like [12] ensure safety by essentially requiring pairwise *syntactic* duality of per-role views at all points in a protocol (called *consistency* [12] or *coherence* [22]). By contrast, our proposed extensions allow additional safe protocols, but also (syntactically) allow protocols with errors that were previously precluded. E.g., consider the choice from `P1` in Sect. 1, where it is safe, but now without the recursion: either `B` or `C` is unsafely left hanging at the end of a session.


```
choice at A { 1() from A to B; } or { 2() from A to C; } [iii]
```

To deal with such additional errors, and those related to explicit connection actions, we validate global types by (1) a lighter set of syntactic conditions, in comparison to standard MPST; complemented by (2) explicit error checking on a *1-bounded model* of the protocol. The key conditions of (1) are:

Role Enabling. For any given choice, we consider the subject (the *at* role) to be *enabled* by default; other roles become enabled after receiving a message. Only enabled roles may connect or send messages to other roles. Role enabling checks that this transitive propagation of the enabled status is respected.

Consistent External Choice Subjects. Every potentially incoming message in an *input* choice (i.e., accepts or receives) must be directed *from* the *same* role.

These basic conditions, in conjunction with the inherent pairing of role-to-role actions in global types, serve the soundness of (2) in the presence of asynchrony and recursion (in general, the state space of an MPST protocol may be unbounded; e.g., P1 in Sect. 1). We note that the latter condition is implicitly imposed by the standard projection in existing MPST [12, 22] (and by projections extended with *merging* [18, 46]), with the additional restriction of directed choice to *send* every *output* choice message to the same role.



We first demonstrate the validation by illustrating some models used by our tool for some previous examples; the details will be covered in Sects. 3 and 4. Initial states are labelled 1; the notation, e.g., $A : B!1$ means A performs the local send $B!1$. (a) is for Ex. [iii]: the two terminals are *unfinished role errors* (Sect. 3.2), where the system is terminated but either B or C is not locally terminated. (b) is for P1 from Sect. 1 assuming output choice fairness, i.e., that both the $B!1$ and $C!2$ options are always viable; this model passes the validation. (c) is the contrary view for P1, where A commits exclusively to a single choice case after the first selection. Our tool additionally constructs this variant to expose such *role progress* violations (Sect. 3.2), where an unfinished role never progresses along some infinite execution, e.g., C does not progress in the cycle between 2 and 3. (a) is not affected by the fairness assumption, as there is no recursion.

The “unfair” model for P2 (not shown) has the same structure as (c), but with connects/disconnects in place of the sends/receives. It would *not* violate progress because either B or C remains in a local connection-accept “guard” state, which is not considered unfinished (rather, “inactive”). **TravelAgency** satisfies progress (i.e., wrt. \mathcal{S}) regardless of output choice fairness for the same reason.

3 MPST with Explicit Connection Actions

3.1 Global Types, Local Types and Sessions

Syntax. A core syntax of *global types* G and *local types* L is defined in Fig. 2. Global types have *guarded choices* $\sum_{i \in I} \pi_i . G_i$, with *connection* $r \rightarrow r' : l$, *messaging* $r \rightarrow r' : l$ and *disconnection* $r \# r'$ actions; *recursion* $\mu X . G$ and X ; and *termination* end . As an example, **TravelAgency** from Fig. 1 may be written (assuming

Roles	$A, B, \dots \in \mathbb{R}$	ranged over by	r, r', \dots
Message labels	$1, 2, 3, \dots \in \mathbb{L}$		l, l', \dots
Recursion variables	$X, Y, \dots \in \mathbb{X}$		X, Y, \dots
Paired interactions	$(\mathbb{R} \times \{\rightarrow, \twoheadrightarrow\} \times \mathbb{R} \times \mathbb{L}) \cup (\mathbb{R} \times \{\#\} \times \mathbb{R})$		π, π', \dots
Localised actions	$\mathbb{A} \subseteq (\mathbb{R} \times \{!, ?, !!, ??\} \times \mathbb{L}) \cup (\mathbb{R} \times \{\#\})$		α, α', \dots

$$G ::= \sum_{i \in I} \pi_i.G_i \mid \mu X.G \mid X \mid \text{end} \quad L ::= \sum_{i \in I} \alpha_i.L_i \mid \mu X.L \mid X \mid \text{end} \quad |I| \geq 1$$

$$r_1 \dagger r_2 : l.G \upharpoonright_{\Delta} r = \begin{cases} r_2 \llbracket \dagger \rrbracket \bullet l.(G \upharpoonright_{\emptyset} r) & r = r_1 & r_1 \# r_2.G \upharpoonright_{\Delta} r = \\ r_1 \llbracket \dagger \rrbracket \circ l.(G \upharpoonright_{\emptyset} r) & r = r_2 & \begin{cases} r' \# \bullet.(G \upharpoonright_{\emptyset} r) & r, r' \in \{r_1, r_2\}, r \neq r' \\ G \upharpoonright_{\Delta} r & \text{otherwise} \end{cases} \\ G \upharpoonright_{\Delta} r & r \notin \{r_1, r_2\} \end{cases}$$

$$\llbracket \rightarrow \rrbracket \bullet = ! \quad \llbracket \rightarrow \rrbracket \circ = ? \quad \llbracket \twoheadrightarrow \rrbracket \bullet = !! \quad \llbracket \twoheadrightarrow \rrbracket \circ = ??$$

$$\sum_{i \in I} G_i \upharpoonright_{\Delta} r = \begin{cases} X \text{ (resp. end)} & \forall i \in I. G_i \upharpoonright_{\Delta} r = X \text{ (resp. } \forall i \in I. G_i \upharpoonright_{\Delta} r = \text{end)} \\ \sum_{j \in J \subseteq I} (L_j = G_j \upharpoonright_{\Delta} r) & |J| > 0, \forall k \in I \setminus J (G_k \upharpoonright_{\Delta} r = \text{end or } X \in \Delta), \text{ and} \\ & \text{either } \begin{cases} \forall j \in J. L_j = \alpha_j^{\bullet}.L'_j \\ \exists r' \forall j \in J. L_j = \alpha_j^{\circ}.L'_j \wedge \text{subj}(\alpha_j^{\circ}) = r' \end{cases} \end{cases}$$

$$\mu X.G \upharpoonright_{\Delta} r = \begin{cases} \text{end} & G \upharpoonright_{\Delta \cup \{X\}} r = X' \text{ or end} & X \upharpoonright_{\Delta} r = X \\ \mu X.(G \upharpoonright_{\Delta \cup \{X\}} r) & \text{otherwise} & \text{end} \upharpoonright_{\Delta} r = \text{end} \end{cases}$$

Fig. 2. Core syntax and global-to-local type projection.

an empty label `nil` for the initial connect, and “flattening” the nested choices):

$C \rightarrow A : \text{nil}.\mu\text{TravelAgency}.$ ($C \rightarrow A : \text{query}.A \rightarrow C : \text{quote}.\text{TravelAgency}$
 $+ C \rightarrow S : \text{pay}.S \rightarrow C : \text{confirm}.C \rightarrow A : \text{acctpt}.\text{end} + C \rightarrow A : \text{reject}.\text{end}$)

Local types are the same except for localised actions: *connect* $r!!l$, *accept* $r??l$, *send* $r!l$, *receive* $r?l$, and *disconnect* $r\#$. For a local action $\alpha = r \dagger l$, the annotation α° means $\dagger \in \{?, ??\}$; and α^{\bullet} means either α with $\dagger \in \{!, !!\}$ or an action $r\#$. We sometimes omit `end`.

The *projection of G onto r* , written $G \upharpoonright r$, is the L given by $G \upharpoonright_{\emptyset} r$ in Fig. 2, where the Δ is a set $\{X_i\}_{i \in I}$. Our projection is more “relaxed” than in standard MPST, in that we seek only to regulate some basic conditions to support the later validation (see below). Δ is simply used to prune X that become unguarded in choices during projection onto r , when the recursive path does not involve r ; e.g., projecting `TravelAgency` onto S : $C??\text{pay}.C!\text{confirm}.\text{end}$). Note: we make certain simplifications for this core formulation, e.g., we omit payload types and flattening of nested choice projections [23].

We assume some basic constraints (typical to MPST) on any given G . **(1)** For all $\pi_{\dagger} = r \dagger r' : l$, $\dagger \in \{\rightarrow, \twoheadrightarrow\}$, and all $\pi_{\#} = r\#r'$, we require $r \neq r'$. We then define: $\text{subj}(\pi_{\dagger}) = \{r\}$, $\text{obj}(\pi_{\dagger}) = \{r'\}$, $\text{lab}(\pi_{\dagger}) = l$; and $\text{subj}(\pi_{\#}) = \{r, r'\}$, $\text{obj}(\pi_{\#}) = \emptyset$. **(2)** G is closed, i.e., has no free recursion variables. **(3)** G features only deterministic choices in its projections. We write: $r \in G$ to mean r occurs in G ; and $\alpha \in L$ to mean $L' = \sum_{i \in I} \alpha_i.L_i$, where L' is obtained from L by some number (possibly zero) of recursion unfoldings, with $\alpha = \alpha_i$ for some i .

$$\begin{array}{c}
 \text{(Sessions)} \quad S ::= (P, Q) \quad P ::= \{L_r\}_{r \in \mathbb{R}} \quad Q : (\mathbb{R} \times \mathbb{R}) \mapsto \{\perp\} \cup \vec{l} \\
 \\
 \text{[CONN]} \quad \frac{\exists i' \in I, j' \in J \quad \alpha_{i'} = r'!!l \quad \alpha_{j'} = r'??l \quad Q(r, r') = Q(r', r) = \perp}{(\{\sum_{i \in I} \alpha_i.L_i\}_r, (\sum_{j \in J} \alpha'_{j'}.L'_{j'})_{r'}) \cup P, Q \rightarrow_k (\{(L_{i'})_r, (L'_{j'})_{r'}\} \cup P, Q[r, r' \mapsto \epsilon][r', r \mapsto \epsilon])} \\
 \\
 \text{[SEND]} \quad \frac{\exists j \in I \quad \alpha_j = r'!l \quad Q(r, r') \neq \perp \quad Q(r', r) = \vec{l} \quad |\vec{l}| < k}{(\{\sum_{i \in I} \alpha_i.L_i\}_r \cup P, Q) \rightarrow_k (\{L_j\}_r \cup P, Q[r', r \mapsto \vec{l}.l])} \\
 \\
 \text{[RECV]} \quad \frac{\exists j \in I \quad \alpha_j = r'?!l \quad Q(r, r') = l.\vec{l}}{(\{\sum_{i \in I} \alpha_i.L_i\}_r \cup P, Q) \rightarrow_k (\{L_j\}_r \cup P, Q[r, r' \mapsto \vec{l}])} \\
 \\
 \text{[DIS]} \quad \frac{Q(r, r') = \epsilon}{(\{r' \# .L\}_r \cup P, Q) \rightarrow_k (\{L\}_r \cup P, Q[r, r' \mapsto \perp])} \quad \text{[REC]} \quad \frac{(\{L[\mu X.L/X]\}_r \cup P, Q) \rightarrow_k (P', Q')}{(\{\mu X.L\}_r \cup P, Q) \rightarrow_k (P', Q')}
 \end{array}$$

Fig. 3. Sessions (pairs of participants and message queues), and session reduction.

(Unfolding is the substitution on recursions $\text{unf}(\mu X.G) = G[\mu X.G/X]$; $\text{unf}(G) = G$ otherwise.) We use \mathbb{R}_G to denote $\{r \mid r \in G\}$, omitting the subscript G where clear from context.

Well-formed global type. For a given G , let $\varphi(G)$ be the global type resulting from the *once-unfolding* of every recursion $\mu X.G'$ occurring within G (defined by $\varphi(\mu X.G) = \varphi(G[\text{end}/X])$, and homomorphic for the other constructors). *Role enabling* (outlined in Sect. 2) on global types $R \vdash G$, $R \subseteq \mathbb{R}$, is defined by $R \vdash \text{end}$ for any R , and:

$$\frac{\text{subj}(\pi) \subseteq R \quad R \cup \text{obj}(\pi) \vdash G}{R \vdash \pi.G} \quad \frac{|I| > 1 \quad \exists r \in R \forall i \in I. \text{subj}(\pi_i) = \{r\} \wedge \{r\} \cup \text{obj}(\pi_i) \vdash G_i}{R \vdash \sum_{i \in I} \pi_i.G_i}$$

A global type G is *well-formed*, $\text{wf}(G)$, if $\mathbb{R}_G \vdash \varphi(G)$, and for all $r \in \mathbb{R}_G$, $G \upharpoonright r$ is *defined*. A consequence is that disconnects are not prefixes in non-unary choices. Also, every local choice in a projection of a $\text{wf}(G)$ comprises only α^\bullet or α° actions, with a consistent subject r in all cases of the latter.

Sessions (Fig. 3) are pairs of a set of *participant* local types P and *inter-role message queues* Q . \perp designates a *disconnected* queue. We use the notation $Q[K \mapsto V]$ to mean Q' where $Q'(K) = V$, and $Q'(K') = Q(K')$ for $K \neq K'$. *Session reduction* (Fig. 3), $S \rightarrow_k S'$, is parameterised on a maximum queue size $k \in \mathbb{N}_1 \cup \{\omega\}$. If two roles are mutually disconnected, [CONN] establishes a connection, synchronising on a common label l . If both sides are connected, [SEND] asynchronously appends a message to destination queue if there is space. If the local queue is still connected: [RECV] consumes the first message, if any; and [DIS] disconnects the queue if it is empty.

For a $\text{wf}(G)$ with roles \mathbb{R} , we define: **(1)** \rightarrow_k^* is the reflexive and transitive closure of \rightarrow_k ; **(2)** the *k-reachable set* of a session S for some k is $RS_k(S) = \{S' \mid S \rightarrow_k^* S'\}$; we say $S' \in RS_k(S)$ is *k-reachable from S*; **(3)** the *initial session* is the session $S_0 = (\{G \upharpoonright r\}_{r \in \mathbb{R}}, Q_{\mathbb{R}_0})$, where $Q_{\mathbb{R}_0} = \{r, r' \mapsto \perp \mid r, r' \in \mathbb{R}\}$; and **(4)** a *k-final session* S is such that $\nexists S'(S \rightarrow_k S')$. We may annotate a reduction

step $S \xrightarrow{r} S'$ by a *subject role* r of the step: in Fig. 3, in [SEND], [RECV] and [DIS] the subject is r ; in [CONN], both r and r' are subjects. Given S , r and k , $S \xrightarrow{r}_k$ stands for $\exists S'(S \xrightarrow{r}_k S')$. For $k = \omega$, we often omit ω .

3.2 MPST Safety and Progress

The following defines MPST safety errors and progress for this formulation. Assume a wf(G) with initial session S_0 and $S \in RS_k(S_0)$ for some k . For $r \in G$, we say: r is *inactive* in S , with $S = (P, Q)$ and $L_r \in P$, if (1) $L_r = \text{end}$; or (2) $L_r = G \upharpoonright r$ and $\forall \alpha \in L_r \exists l(\alpha = r'??l)$. Otherwise, r is *active* in S .

Then, session $S = (P, Q)$ is a *safety error*, Err, if:

(i) $L_r \in P$ and one of the following holds:

- (Connection error) $\alpha \in L_r$, $\alpha = r'!!l$ or $r'??l$, and $Q(r, r') \neq \perp$;
- (Disconnect error) $r' \# \in L_r$ and $Q(r, r') \neq \epsilon$;
- (Unconnected error) $\alpha \in L_r$, $\alpha = r'!l$ or $r'?!l$, and $Q(r, r') = \perp$;
- (Synchronisation error) $r'!!l \in L_r$, $(\sum_{i \in I} r'??l_i.L_i)_{r'} \in P$, and $l \notin \{l_i\}_{i \in I}$;
- (Reception error) $L_r = \sum_{i \in I} r'?!l_i.L_i$, $Q(r, r') = l \cdot \vec{l}$ and $l \notin \{l_i\}_{i \in I}$;

or (ii) S is either:

- (Unfinished role) S is k -final and $r \in G$ is active in S ;
- (Orphan message) $r \in G$ is inactive in S and $\exists r'(Q(r, r') \notin \{\perp, \epsilon\})$.

Session S satisfies *k -progress* if, for all $S' = (P, Q) \in RS_k(S)$, we have both: (*Role progress*) for all $r \in \mathbb{R}$, if r is active in S' , then $S' \xrightarrow{r}_k^* S''$; and (*Eventual reception*) if $Q(r, r') = l \cdot \vec{l}$, then $S' \xrightarrow{r}_k^* (P', Q')$ where $L_r \in P'$ and $r'?!l \in L_r$.

A session S is *k -safe* if $\nexists \text{Err} \in RS_k(S)$. We simply say session S is *safe* if it is ω -safe; and S satisfies *progress* if it satisfies ω -progress.

The following establishes the soundness of our framework. Our approach is to adapt the CFSM-based methodology of [6, 18], by reworking the notion of *multiparty compatibility* developed there, in terms of the syntactic conditions and explicitly checked 1-bounded properties of our extended setting. The omitted/expanded definitions and proofs can be found at [23].

Theorem 1 (*Soundness of 1-bounded validation*). *Let S_0 be the initial session of a wf(G) that is 1-safe and satisfies 1-progress. Then S_0 is safe and satisfies progress.*

4 Implementation


4.1 Modelling MPSTs by CFSMs with Dynamic Connections

We present a prototype implementation [43] that adapts the preceding formulation by constructing and checking explicit state models of our extended global types, based on a correspondence between MPST and communicating FSMs

(CFSMs) [15,18,30]. In this setting, our extensions correspond to CFSMs with *dynamic connection actions*. An *Endpoint FSM* (EFSM) for a role is:

$$(EFSM) M = (\mathbb{S}, \mathbb{R}, s_0, \mathbb{L}, \delta) \quad (\text{States}) s, s', \dots \in \mathbb{S} \quad (\text{Transitions}) \delta \subseteq \mathbb{S} \times \mathbb{A} \times \mathbb{S}$$

where s_0 is the *initial state*; \mathbb{R}, \mathbb{L} and \mathbb{A} are as defined in Fig. 2. We write $\delta(s)$ to denote $\{\alpha \mid \exists s'. \delta(s, \alpha) = s'\}$. EFSMs are given by a (straightforward) translation from local types, for which we omit the full details [23]: an EFSM essentially captures the structure of the syntactic local type with recursions reflected as cycles. E.g., for \mathbf{C} in `TravelAgency` (Fig. 1), omitting payload types:



The execution of EFSM systems is adapted from basic CFSMs [9] following Fig. 3 in the expected way [23]. Then, assuming an *initial configuration* c_0 (the system with all endpoints in their initial EFSM states and unconnected) for a $wf(G)$, the (base) *model of G* is the set of configurations that can be reached by *1-bounded* execution from c_0 . We remark that the model of a $wf(G)$ is finite.

Based on Sect. 3.2, G can be validated by its model as follows. The MPST safety errors pertain to individual configurations: this allows to simply check each configuration by adapting the Err-cases to this setting. E.g., an *unfinished role* error is a terminal configuration where role r is in a non-terminal state s_r , and s_r is not an accept-guarded initial state. MPST progress for potentially non-terminating sessions can be characterised on the finite model in terms of closed subsets of mutually reachable configurations (sometimes called *terminal sets*). E.g., a *role progress* violation manifests as such a closure in which an active role is not involved in any transition (e.g., configs. 2 and 3, wrt. \mathbf{C} , in (c) on p. 6).


Choice subtyping vs. progress. A projected local choice is either an output choice (connects, sends) or an input choice (accepts, receives). While input choices are driven by the received message, output choices are driven by *process-level* procedures that global and local types abstract from. The notion of *session subtyping* [13,20] was developed to allow more flexible implementations against a local type. E.g., the projection of $\mathbf{P1}$ from Sect. 1 for \mathbf{A} is $\mu X. (\mathbf{B}!1.X + \mathbf{C}!2.X)$ which says \mathbf{A} repeatedly has the choice of sending 1 to \mathbf{B} or 2 to \mathbf{C} : intuitively, it is *safe* here to implement an \mathbf{A} that always opts to send 1 (e.g., a process $P(x) = x \oplus (\mathbf{B}, 1).P\langle x \rangle$, where x is \mathbf{A} 's session channel, \oplus is the select primitive [12]). For our relaxed form of multiparty choice, however, such a (naive) interpretation of subtyping raises the possibility of *progress* errors (in this case, for \mathbf{C}).

To allow our validation approach to be integrated with the various methods of verifying local types in real-world languages, we consider this issue from the perspective of two basic assumptions on implementations of output choices. One is to simply assume *output choice fairness* (the basic interpretation that an infinite execution of an output choice selects each recursive case infinitely many times), which corresponds to the model construction defined so far.

The other interpretation is developed as a “worst case” view, where we do not assume any direct support for session typing or subtyping (fair or otherwise) in the target language (e.g., native Java), and allow the implementation of every recursive output choice to be reduced to only ever following one particular case. Our tool implements this notion as a transformation on each EFSM, by refining the continuations of output choices such that the *same* case is always selected if that choice is repeated in the future. We outline the transformation below (see [23] for the definition):

- For each non-unary output choice s^\bullet , we clone the subgraph reachable via an action $\alpha \in \delta(s^\bullet)$ in each case that s^\bullet is reachable via α , i.e., if $s^\bullet \in RS(\delta(s^\bullet, \alpha))$.
- In each subgraph cloned via α , all $\alpha' \in \delta(s^\bullet)$ edges, s.t. $\alpha' \neq \alpha$, are pruned from the clone of s^\bullet . We redirect the α -edge from s^\bullet to the clone of its successor $\delta(s^\bullet, \alpha)$ in the cloned subgraph. (States no longer connected are discarded.)
- This transformation is applied recursively on the cloned subgraphs, until every recursive output choice is reduced to a single action.

This transformation reflects endpoint implementations that push output choice subtyping to exercise a *minimum* number of different recursive cases along a path. To expose progress violations under subtyping when fairness is not assumed, our tool uses the transformed EFSMs to additionally construct and check the “unfair” 1-bounded global model in the same manner as above.



We illustrate some examples. (d) is the base EFSM, i.e., assuming output choice fairness, for A in P1 from Sect. 1. (e) is the transformed EFSM: if A starts by selecting the 1 case it will continue to select this case only; similarly for 2. (The transformation does not change B or C.) Using (e) gives the global model for P1 in (c) on p. 6, raising the role progress violations for B and C. By contrast, (f) is the transformed EFSM for A in P2 from Sect. 1: as in (e), A commits exclusively to whichever case is selected first. However, P2 does not violate progress, despite the transformation of A in (f), because the involvement of C is guarded by the initial connection-accept actions in (g); similarly for B.

4.2 Type-Checking Endpoint Programs by Local Type Projections

Java endpoint implementation via API generation. We demonstrate an integration of the above developments with an existing approach for using local types to verify endpoint programs. Concretely, we extend the approach of [24], to *generate* Java APIs for implementing each role of a global type, including explicit connection actions, via the translation of projections to EFSMs. The idea is to reify each EFSM state as a Java class for a *state-specific* channel, offering methods for exactly the permitted I/O actions. These channel classes are linked by

```

1  TravelAgency sess = new TravelAgency(); // Generated session class
2  try (ExplicitEndpoint<TravelAgency, C> ep = new ExplicitEndpoint<>(sess, C) {
3      Buf<Integer> b = new Buf<>();
4      TravelAgency_C_2 C2 = new TravelAgency_C_1(ep) // Generated channel classes
5          .connect(A, SocketChannelEndpoint::new, host_A, port_A); // TCP client
6      for (int i = 0; i < queries.length; i++) // Assume queries: String[]
7          C2 = C2.send(A, query, queries[i]).receive(A, quote, b);
8      C2.connect(S, SocketChannelEndpoint::new, host_S, port_S, // TCP client
9          pay, "..paymentInfo..").receive(S, confirm, b)
10     .send(A, accpt, b.val); // C simplified to always accept the quote
11 } // (reject option unused)

```

Fig. 4. Safe Java implementation of C in `TravelAgency` (Fig. 1) using generated APIs.

setting the return type of each method to its successor state. Session safety is assured by static (Java) typing of the I/O method calls, combined with run-time checks (built into the API) that each instance of a channel class is used *exactly once*, for the linear aspect of session typing. An endpoint implementation thus proceeds, from a channel instance of the initial state, by calling one I/O method on the current channel to obtain the next, up to the end of the session (if any).

Figure 4 illustrates the incorporation of explicit connect, accept and disconnect actions from local types into the API generated for C in `TravelAgency`; this code can be compared against the EFSM on p. 10. `TravelAgency_C_1` is the initial state channel (cf. EFSM state 1), for which the *only* permitted I/O method is the `connect` to A ; attempting any other session operation is simply a Java type error. (The various constants, such as A and `query`, are *singleton type values* in the API.) The `connect` returns a new instance of `TravelAgency_C_2`, offering exactly the mixed choice between the non-blocking `query` (line 7) or `reject` (*unused*, cf. Sect. 4.1, output choice subtyping) to A , or the blocking `connect` to S (line 8).

If the programmer respects the linear channel usage condition of the generated API, as in Fig. 4, then Java typing statically ensures the session code (I/O actions and message types) follows its local type. The only way to violate the protocol is to violate linearity, in which case the API will raise an exception without actually performing the offending I/O action. Our toolchain, from validated global types to generated APIs, thus assures safe executions of endpoint implementations up to premature termination.

Correlating dynamic binary connections in multiparty sessions. Even aside from explicit connections, session initiation is one aspect in which applications of session type theory, binary and multiparty, to real distributed systems raises some implementation issues. The standard π -calculus theory assumes a so-called *shared channel* used by all the participants for the initiation synchronisation.² The formal typing checks, on a “centralised” view of the entire process system, that each and every role is played by a compliant process, initiated via the shared channel. These assumptions transfer to our distributed, binary-connection programs as relying on correct host and port argument values in, e.g., the

² E.g., a in $a[1](y).P_1 \mid \dots \mid a[n-1](y).P_{n-1} \mid \bar{a}[n](y).P_n$, initiating a session between n processes [12].

`connect` calls in `C` in Fig. 4 (lines 5 and 8); similarly for the arguments to the `SocketChannelServer` constructor and `accept` call in the `A` and `S` programs [23].

Existing π -calculus systems could be naively adapted to explicit connection actions by assigning a (binary) shared channel to each accept-point in the session, since the type for any given point in a protocol is fixed. Unfortunately, reusing a shared channel for dynamic accepts across *concurrent* sessions may lead to incorrect *correlation* of the underlying binary connections. E.g., consider $A \rightarrow B..A \rightarrow C..B \rightarrow C..$, where the C process uses multithreading to concurrently serve multiple sessions: if the *same* shared channel is used to accept *all* connections from the A 's, and likewise for B 's, there is no inherent guarantee that the connection accepted from a B by a given server thread will belong to the same *session* as the earlier connection from A , despite being of the expected type.

In practice, the correlation of connections to sessions may be handled by various mechanisms, such as passing session identifiers or port values. Consider the version of the `Pay` subprotocol (from Fig. 1), modified to use port passing (cf. FTP [39]), on the left:

<pre>choice at C { accept() from C to A; connect A to S; port(Int) from S to A; port(Int) from A to C; pay(Str) connect C to S; confirm(Int) from S to C; } or { reject() from C to A; }</pre>	<pre>// Extended Scribble annotations: [iv] // S opens a (fresh) Int port for C port(p:Int) from S to A; @open=p:C port(p) from A to C; // A forwards p pay(Str) connect C to S; @port=p // C connects using p as the port</pre>
--	--

`C` sends `accept` to `A`, and then `A` connects to `S`; `S` sends `A` an `Int` port value, which `A` forwards to `C`; `C` then connects to `S` at that port. To capture this intent explicitly, we adapt an extension of Scribble with assertions [34] to support the specification on the right. In general, *value*-based constraints, like forwarding and connecting to `p`, can be generated into the API as implicit run-time Java assertions. However, we take advantage of the API generation approach to directly generate *statically* safe operations for these actions. *N.B.*, in the following, *port* is simply the message label API constant; assigning, sending and using the actual port value is safely handled *internally* by the generated operations.

```
In S: s.send(A, port).accept(C, pay, b).. // 'port' is the msg label (API const.)
// API internally opens fresh port p, and sends value; accepts conn. on p
In A: A.receive(S, port).send(C, port)..
// API internally caches the received value of p; and forwards that value
In C: C.receive(A, port).connect(S, SocketChanEndpoint::new, host_S, pay, "..")..
// API internally caches the value of p; and connects using p as the port
```

This combination of explicit connection actions, assertions, and typed API generation is essentially a practical realisation of (private) *shared channel passing* from session π -calculi for our binary connection setting in Java.

To facilitate integration with some existing implementations of session typed languages, our toolchain also supports an optional syntactic restriction on types where: each projection of a Scribble protocol may contain at most one *accept-choice* constructor, and only as the top-most choice constructor (cf. the commonly used replicated-server process primitives in process calculus works).

This constraint allows many useful explicit connection action patterns, including nested connects and recursive accepts, while ruling out correlation errors; apart from Ex. [iv], all of the examples in this paper satisfy this constraint.

5 Related Work and Concluding Remarks

Dynamic participants in typed process calculi and message sequence charts. To our knowledge, this paper is the first session types work that allows a single session to have optional roles, and dynamic joining and leaving of roles.

[16] presents a version of session types where a role designates a dynamic *set* of one or more participant processes. Their system does not support optional nor dynamic roles (every role is played by at least one process; the number of processes varies, but the set of *active roles* is fixed). It relies on a special-purpose run-time locking mechanism to block dynamically joining participants until some safe entry point, hindering its use in existing applications. Implementations of sessions in Python [15] and Erlang [19] have used a notion of *subsession* [14] as a coarse-grained mechanism for dynamically introducing participants. The idea is to launch a *separate* child session, by the heavyweight atomic multiparty initiation, involving a subset of the current participants along with other new participants; unlike this paper, where additional roles enter the same, running session by the connect and accept actions between the two relevant participants.

The *conversation calculus* [10] models conversations between dynamic *contexts*. A behavioural typing ensures error-freedom by characterising processes more directly; types do not relate to roles, as in MPST. Their notion of dynamic joining is more abstract (akin to standard MPST initiation), allowing a context n to interact with all other conversation members after a single atomic join action by n ; whereas our explicit communication actions are designed to map more closely to concrete operations in standard network APIs.

Dynamic message sequence charts (DMSCs) in [31] support fork-join patterns with potentially unbounded processes. Model checking against a monadic second order logic is decidable, but temporal properties are not studied. [7] studies the implementability of *dynamic communication automata* (DCA) [8] against MSCs as specifications. The focus of study of DCA and DMSCs is more about dynamic *process spawning*; whereas we target dynamic *connections* (and disconnects) between a set of roles with specific concern for MPST safety and progress. Our implementation goes another “level” down from the automata model, applying the validated session types to Java programs with consideration of issues such as choice subtyping and connection correlation.

Well-formedness of session types and choreographies. Various techniques involving bounded executions have been used for multiparty protocols and choreographies. [3, 4, 41] positions choreography realisability in terms of *synchronisability*, an equivalence between a synchronous global model and the 1-bounded execution of local FSMs; this reflects a stricter perspective of protocol compliance, demanding stronger causality between global steps than session type safety. Their communication model has a single input queue per endpoint, while asynchronous

session types has a separate input queue per peer: certain patterns are not synchronisable in the former while valid in the latter. [2] develops more general realisability conditions (in the single-queue model) than the above by determining an upper-bound on queue sizes wrt. equivalent behaviours. Our validation of MPST with explicit connection actions remains within a 1-bounded model.

[18] characterises standard MPST wrt. CFSMs by *multiparty compatibility*, a well-formedness condition expressed in terms of 1-reachability; it corresponds to the syntactic restrictions of standard MPST in ensuring safety. This paper relaxes some of these restrictions with other extensions, by our 1-bounded validation, to support our use cases. [30] develops a bottom-up synthesis of graphical choreographies from CFSMs via a correspondence between synchronous global models and local CFSMs. These works and the above works on choreographies: (1) do not support patterns with optional or dynamic participants; and (2) study single, pre-connected sessions in isolation without consideration of certain issues of implementing endpoint programs in practice (type checking, subtyping, concurrent connection correlation).

Advanced subtyping of local types with respect to liveness is studied theoretically in [37]. Our present work is based on a coarser-grained treatment of fairness in the global model, to cater for applications to existing (mainstream) languages where it may be difficult to precisely enforce a particular subtyping for sessions via the native type system. We plan to investigate the potential for incorporating their techniques into our approach in future work.

Implementations of session types. The existing version of Scribble [24, 47] follows the established theory through syntactic restrictions to ensure safety (e.g., the same set of roles must be involved in every choice case, precluding optional participation). [24] concerns only the use of local types for API generation; it has no formalism, and does not discuss global type validation or projection. This paper is motivated by use cases to relax existing restrictions and add explicit connection actions to types. [38] develops a tool for checking or testing compatibility, adapted from [18], in a local system of abstract concurrent objects. It does not consider global types nor endpoint programs.

Recent implementation works [24–27, 32, 35, 36, 40, 42, 45], as discussed in Sect. 1, have focused more on applying standard session types, rather than developing session types to better support real use cases. This is in contrast to the range of primarily theoretical extensions (e.g., time [6, 33], asynchronous interrupts [15], nested subsessions [14], assertions [5], role parameterisation [46], event handling [29], multi-process roles [16], etc.), which complicates tool implementation because each has its own specific restrictions to treat the subtleties of its setting. The approach of this paper, shifting the emphasis from outright syntactic well-formedness to a more uniform validation of the types, may be one way to help bring some of these scattered features (and those in this paper) together in practical MPST implementations. We plan to investigate such directions in future work, in addition to closer integrations of MPST tools, that treat concepts like role projections, endpoint program typing, subtyping and channel passing, with established model checking tools and optimisations.

Acknowledgements. We thank Gary Brown and Steve Ross-Talbot for collaborations, and Romyana Neykova for comments. This work is partially supported by EPSRC projects EP/K034413/1, EP/K011715/1, EP/L00058X/1, EP/N027833/1, and EP/N028201/1; and by EU FP7 612985 (UPSCALE).

References

1. Ancona, D., et al.: Behavioral types in programming languages. *Found. Trends Program. Lang.* **3**(2–3), 95–230 (2016)
2. Basu, S., Bultan, T.: Automatic verification of interactions in asynchronous systems with unbounded buffers. In: ASE 2014, pp. 743–754. ACM (2014)
3. Basu, S., Bultan, T.: Automated choreography repair. In: Stevens, P., Wařowski, A. (eds.) FASE 2016. LNCS, vol. 9633, pp. 13–30. Springer, Heidelberg (2016). doi:[10.1007/978-3-662-49665-7_2](https://doi.org/10.1007/978-3-662-49665-7_2)
4. Basu, S., Bultan, T., Ouederni, M.: Deciding choreography realizability. In: POPL 2012, pp. 191–202. ACM (2012)
5. Bocchi, L., Honda, K., Tuosto, E., Yoshida, N.: A theory of design-by-contract for distributed multiparty interactions. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 162–176. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15375-4_12](https://doi.org/10.1007/978-3-642-15375-4_12)
6. Bocchi, L., Lange, J., Yoshida, N.: Meeting deadlines together. In: CONCUR 2015. LIPIcs, vol. 42, pp. 283–296. Schloss Dagstuhl (2015)
7. Bollig, B., Cyriac, A., H elou et, L., Kara, A., Schwentick, T.: Dynamic communicating automata and branching high-level MSCs. In: Dediu, A.-H., Mart ın-Vide, C., Truthe, B. (eds.) LATA 2013. LNCS, vol. 7810, pp. 177–189. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-37064-9_17](https://doi.org/10.1007/978-3-642-37064-9_17)
8. Bollig, B., H elou et, L.: Realizability of dynamic MSC languages. In: Ablayev, F., Mayr, E.W. (eds.) CSR 2010. LNCS, vol. 6072, pp. 48–59. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-13182-0_5](https://doi.org/10.1007/978-3-642-13182-0_5)
9. Brand, D., Zafiropulo, P.: On communicating finite-state machines. *J. ACM* **30**, 323–342 (1983)
10. Caires, L., Vieira, H.T.: Conversation types. *Theor. Comput. Sci.* **411**(51–52), 4399–4440 (2010)
11. C ec e, G., Finkel, A.: Verification of programs with half-duplex communication. *Inf. Comput.* **202**(2), 166–190 (2005)
12. Coppo, M., Dezani-Ciancaglini, M., Yoshida, N., Padovani, L.: Global progress for dynamically interleaved multiparty sessions. *Math. Struct. Comput. Sci.* **760**, 1–65 (2015)
13. Demangeon, R., Honda, K.: Full abstraction in a subtyped pi-calculus with linear types. In: Katoen, J.-P., K onig, B. (eds.) CONCUR 2011. LNCS, vol. 6901, pp. 280–296. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-23217-6_19](https://doi.org/10.1007/978-3-642-23217-6_19)
14. Demangeon, R., Honda, K.: Nested protocols in session types. In: Koutny, M., Ulidowski, I. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 272–286. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-32940-1_20](https://doi.org/10.1007/978-3-642-32940-1_20)
15. Demangeon, R., Honda, K., Hu, R., Neykova, R., Yoshida, N.: Practical interruptible conversations: distributed dynamic verification with multiparty session types and python. In: Formal Methods in System Design, pp. 1–29 (2015)
16. Deni elou, P.-M., Yoshida, N.: Dynamic multirole session types. In: POPL 2011, pp. 435–446. ACM (2011)

17. Deniérou, P.-M., Yoshida, N.: Multiparty session types meet communicating automata. In: Seidl, H. (ed.) ESOP 2012. LNCS, vol. 7211, pp. 194–213. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-28869-2_10](https://doi.org/10.1007/978-3-642-28869-2_10)
18. Deniérou, P.-M., Yoshida, N.: Multiparty compatibility in communicating automata: characterisation and synthesis of global session types. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) ICALP 2013. LNCS, vol. 7966, pp. 174–186. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-39212-2_18](https://doi.org/10.1007/978-3-642-39212-2_18)
19. Fowler, S.: An erlang implementation of multiparty session actors. In: ICE 2016. EPTCS, vol. 223, pp. 36–50 (2016)
20. Gay, S., Hole, M.: Subtyping for session types in the pi-calculus. *Acta Informatica* **42**(2/3), 191–225 (2005)
21. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: POPL 2008, pp. 273–284. ACM (2008)
22. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. *J. ACM* **63**(1), 9 (2016)
23. Hu, R., Yoshida, N.: Explicit Connection Actions in Multiparty Session Types (Long Version). <https://www.doc.ic.ac.uk/rhu/~scribble/explicit.html>
24. Hu, R., Yoshida, N.: Hybrid session verification through endpoint API generation. In: Stevens, P., Wasowski, A. (eds.) FASE 2016. LNCS, vol. 9633, pp. 401–418. Springer, Heidelberg (2016). doi:[10.1007/978-3-662-49665-7_24](https://doi.org/10.1007/978-3-662-49665-7_24)
25. Hu, R., Yoshida, N., Honda, K.: Session-based distributed programming in java. In: Vitek, J. (ed.) ECOOP 2008. LNCS, vol. 5142, pp. 516–541. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70592-5_22](https://doi.org/10.1007/978-3-540-70592-5_22)
26. Imai, K., Yuen, S., Agusa, K.: Session type inference in haskell. In: PLACES. EPTCS, vol. 69, pp. 74–91 (2010)
27. Jespersen, T.B.L., Munksgaard, P., Larsen, K.F.: Session types for rust. In: WGP 2015, pp. 13–22. ACM (2015)
28. Klensin, J.: IETF RFC 5321 Simple Mail Transfer Protocol. <https://tools.ietf.org/html/rfc5321>
29. Kouzapas, D., Yoshida, N., Hu, R., Honda, K.: On asynchronous eventful session semantics. *Math. Struct. Comput. Sci.* **26**(2), 303–364 (2016)
30. Lange, J., Tuosto, E., Yoshida, N.: From communicating machines to graphical choreographies. In: POPL 2015, pp. 221–232. ACM (2015)
31. Leucker, M., Madhusudan, P., Mukhopadhyay, S.: Dynamic message sequence charts. In: Agrawal, M., Seth, A. (eds.) FSTTCS 2002. LNCS, vol. 2556, pp. 253–264. Springer, Heidelberg (2002). doi:[10.1007/3-540-36206-1_23](https://doi.org/10.1007/3-540-36206-1_23)
32. Lindley, S., Morris, J.G.: Lightweight Functional Session Types. <http://homepages.inf.ed.ac.uk/slindley/papers/fst-draft-february2015.pdf>
33. Neykova, R., Bocchi, L., Yoshida, N.: Timed runtime monitoring for multiparty conversations. In: BEAT 2014. EPTCS, vol. 162, pp. 19–26 (2014)
34. Neykova, R., Yoshida, N., Hu, R.: SPY: local verification of global protocols. In: Legay, A., Bensalem, S. (eds.) RV 2013. LNCS, vol. 8174, pp. 358–363. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-40787-1_25](https://doi.org/10.1007/978-3-642-40787-1_25)
35. Ng, N., Figueiredo Coutinho, J.G., Yoshida, N.: Protocols by default. In: Franke, B. (ed.) CC 2015. LNCS, vol. 9031, pp. 212–232. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-46663-6_11](https://doi.org/10.1007/978-3-662-46663-6_11)
36. Padovani, L.: FuSe homepage. <http://www.di.unito.it/padovani/Software/FuSe/FuSe.html>
37. Padovani, L.: Fair subtyping for multi-party session types. *Math. Struct. Comput. Sci.* **26**(3), 424–464 (2016)

38. Perera, R., Lange, J., Gay, S.J.: Multiparty compatibility for concurrent objects. In: PLACES 2016. EPTCS, vol. 211, pp. 73–82 (2016)
39. Postel, J., Reynolds, J.: IETF RFC 959 File Transfer Protocol. <https://tools.ietf.org/html/rfc959>
40. Pucella, R., Tov, J.A.: Haskell session types with (almost) no class. In: Haskell 2008, pp. 25–36. ACM (2008)
41. Salaün, G., Bultan, T., Roohi, N.: Realizability of choreographies using process algebra encodings. *IEEE Trans. Serv. Comput.* **5**(3), 290–304 (2012)
42. Scalas, A., Yoshida, N.: Lightweight session programming in scala. In: ECOOP 2016. LIPICs, vol. 56, pp. 21:1–21:28. Schloss Dagstuhl (2016)
43. Scribble.: GitHub repository. <https://github.com/scribble/scribble-java>
44. Scribble homepage. <http://www.scribble.org>
45. Sivaramakrishnan, K.C., Qudeisat, M., Ziarek, L., Nagaraj, K., Eugster, P.: Efficient sessions. *Sci. Comput. Program.* **78**(2), 147–167 (2013)
46. Yoshida, N., Deniérou, P.-M., Bejleri, A., Hu, R.: Parameterised multiparty session types. In: Ong, L. (ed.) FoSSaCS 2010. LNCS, vol. 6014, pp. 128–145. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-12032-9_10](https://doi.org/10.1007/978-3-642-12032-9_10)
47. Yoshida, N., Hu, R., Neykova, R., Ng, N.: The scribble protocol language. In: Abadi, M., Lluch Lafuente, A. (eds.) TGC 2013. LNCS, vol. 8358, pp. 22–41. Springer, Cham (2014). doi:[10.1007/978-3-319-05119-2_3](https://doi.org/10.1007/978-3-319-05119-2_3)