

A multiparty multi-session logic

Laura Bocchi¹, Romain Demangeon², and Nobuko Yoshida³

¹University of Leicester, ²Queen Mary, University of London, ³Imperial College London

Abstract. Recent work on the enhancement of typing techniques for multiparty sessions with logical annotations enables, not only the validation of structural properties of the conversations and on the sorts of the messages, but also properties on the actual values exchanged. However, specification and verification of mutual effects of multiple cross-session interactions are still an open problem. We introduce a multiparty logical proof system with virtual states that enables the tractable specification and validation of fine-grained inter-session correctness properties of processes participating in several interleaved sessions. We present a sound and relatively complete static verification method.

1 Introduction

In extensively distributed computing environments, application scenarios often centre around structured conversations among multiple distributed participants. A fundamental challenge is to establish an effective specification and verification method to ensure safety in distributed software, where correctness depends on the state of individual participants and span over multiple conversations and applications. This requirement emerged from our ongoing collaboration with the Ocean Observation Initiative OOI [21], an NSF program to develop a long-term computational infrastructure for environmental ocean observation. The principals within the OOI infrastructure perform interactive activities involving distributed resources, e.g., remote instruments, off-shore sensors, data. It is important to: (1) ensure that the principals carry out each activity (session) in a way that conform a well-defined protocol, (2) express properties that span the single activities (e.g., associate each principal with a credit for resource usage, and ensure that this will always be non negative across sessions¹).

A promising direction is the logical elaboration of types for programming languages [16]. Types offer a stable linkage between the fundamental dynamics of programs and their mathematical abstractions, serving as a highly effective basis for safety assurance. In the context of process algebras, approaches like [5, 13, 18] allow tractable² (e.g., with respect to model checking techniques) validation of properties such as session fidelity, progress, and error freedom. Furthermore, they enable the specification of *global* properties of multiparty interactions, yet enabling modular *local* verification of each principal. The key idea is that conversations are built as the composition of units of design called *sessions* which are specified from a global perspective (i.e., a global session type). Each global type is then *projected*, making the responsibilities of each

¹ This example is taken from the OOI Instrument Control case study and is illustrated in the online report [4].

² In [13, 18] verification is decidable and has linear complexity.

endpoint explicit. Validation guarantees that when each endpoint conforms to its projected specification(s), the resulting conversation conforms to the corresponding global specification(s).

These approaches require to build applications starting from a set of global types that have to be agreed upon by the principals in the network. This assumption, which poses some limitations to the flexibility with which the single local processes are modelled, is reasonable in many scenarios, provided that local processes can be built as the composition of multiple, possibly interleaved, types of sessions. However, one limitation of these approaches is that the properties that they verify are confined to the single multiparty sessions and do not treat stateful specifications incorporating mutual effects of multiple sessions run by a principal.

This paper presents a simple but powerful extension of multiparty session specifications, by enriching the assertion language studied in [5] with capability to refer to *virtual states* local to each network principal. The resulting protocol specifications are called *multiparty stateful assertions (MPSAs)*, and model the skeletal structure of the interactions of a session, the constraints on the exchanged messages and on the branches to be followed, and the *effects* of each interaction on the virtual state. We use *invariants* to express properties, on the state of each principal, that must hold even when several sessions are executed in parallel. Principals in a network hence serve as units of verification: static validation ensures that principals behave as prescribed by *MPSAs* and their invariants are satisfied.

To see the kind of properties we are interested in, consider the following fragment of specification for the dialogue between a ticket allocation server (S) and its client (C), where the server allocates numbered tickets of increasing value to each client in consecutive, *separate* sessions:

$$S \rightarrow C : (y : \text{int}) \{y = S.x\} \langle S.x++ \rangle$$

The protocol between the server and each client is the single message-passing action where S sends C a message of type `int`. The description of this simple distributed application implies behavioural constraints of greater depth than the basic communication actions. The (sender-side) *predicate and effect* for the interaction step, $\{y = S.x\} \langle S.x++ \rangle$, asserts that the message y sent to each client must equal the current value of $S.x$, a state variable x allocated to the *principal* serving as S; and that the local effect of this message send is to increment $S.x$. In this way, S is specified to send incremental values across *consecutive* sessions.

The behaviour described above cannot be encoded by only using the primitives in [5]. In fact, in order to ensure inter-session properties one must discipline concurrent state updates with some mechanism of lock/unlock or atomic access/update, but lock/unlock and atomic access/update can only be described as properties that span over multiple sessions.

Contribution We present a sound and relatively complete validation method for *MPSAs*, based on statically-verifiable proof rules. The most distinctive feature with respect to [5] is the possibility of expressing properties that span several sessions. The decidability/complexity of verification depends on the decidability/complexity of predicate evaluation in the logic that is chosen to express constraints and invariants (Proposition 10). We prove that our analysis is sound (Theorem 13) and complete (Theorem 14)

w.r.t. the semantical satisfaction relation induced by the two labelled transition systems for processes and specifications.

Synopsis In § 2 we present *MPSAs*, that is the language for (stateful) protocol specifications, and their consistency criteria (i.e., well-assertedness). In § 3 we present the calculus for networks of principals, each running a process. In § 4 we give the validation rules of networks against *MPSAs*; their properties are presented in § 5. Related work is discussed in § 6. Use cases from [21], auxiliary definitions, and full proofs can be found in the online report [4].

2 Multiparty assertions with virtual states

In the proposed framework, applications are built as the composition of units of design called *sessions*. Each type of session is specified as a *MPSA*, that is an abstract description of the interactions of the roles of a multiparty session.

The syntax of *MPSAs* is given in Figure 1. *Global assertions* ($\mathcal{G}, \mathcal{G}', \dots$) describe a multiparty session from a global perspective; and *local assertions* ($\mathcal{L}, \mathcal{L}', \dots$) describe it from the perspective of one role. U are types of the message contents, which can be sorts S or local assertions $\langle \mathcal{L} \rangle$ (i.e., for delegation).

$A ::= \text{true} \mid \text{false} \mid e_1 = e_2 \mid \neg A \mid A_1 \wedge A_2 \mid \exists x.A,$	$U ::= S \mid \langle \mathcal{L} \rangle,$	$S ::= \text{bool} \mid \text{int} \mid \dots$
$\mathcal{G} ::= p \rightarrow q : \{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{G}_i\}_{i \in I}$	(G-int) $\mathcal{L} ::= p!\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I}$	(L-sel)
$\mathcal{G}_1 \mid \mathcal{G}_2$	(G-par) $p?\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I}$	(L-bra)
$\mu t\langle y : A' \rangle(x : S)\{A\}.\mathcal{G}$	(G-def) $\mu t\langle y : A' \rangle(x : S)\{A\}.\mathcal{L}$	(L-def)
$t\langle y : A' \rangle$	(G-call) $t\langle y : A' \rangle$	(L-call)
end	(G-end) end	(L-end)

Fig. 1. Global and local MPSAs

For expressing constraints we use *predicates* (A, A', \dots) with the syntax illustrated in Figure 1, although we may use other predicates than equality in examples. Predicates are defined on expressions (e, \dots) containing occurrences of *interaction variables* (x, y, \dots), modelling the content of a message exchanged by the roles in the session, and of *state variables* ($C.x, S.y, \dots$), which are variables of the virtual state local to one role.

Global Assertions Interaction **(G-int)**, $p \rightarrow q : \{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{G}_i\}_{i \in I}$, models a message exchange where role p sends q one of the branch labels l_i and an interaction variable x_i , with x_i binding its occurrences in A_i, E_i , and \mathcal{G}_i . A_i is the predicate which needs to hold for p to select l_i , and which may constrain the values to be sent for x_i . Note that A_i is at the same time an assumption for the receiver q and a constraint for the sender p (i.e., if A_i is violated then the blame is on p). E_i is the update prescribed on the virtual states of p and q , modelling the persistent effects (i.e., with respect to the lifetime of the single session) of that interaction. An update is a vector of assignments

of the form $x := e$, where x is updated by the result of evaluating e in the current state. We assume E does not contain two assignments to the same state variable, and is an atomic action. Assertion **(G-par)** is for parallel composition. The recursive definition **(G-def)** is the guarded recursion definition and defines a recursion parameter x initially set equal to a value satisfying the initialisation predicate A' , with A being an invariant predicate. Global assertions are unfolded implicitly, following an equi-recursive view on types. **(G-call)** is the recursive instantiation and **(G-end)** is the termination.

Hereafter we omit true predicates, empty vectors of variables/updates, and labels of single branches.

Example 1. Consider a session with two roles, C and S. C makes an offer x to S for buying a ticket; S either accepts or refuses the offer. In the former case C spends x credits and receives a ticket, and S earns x credits. Tickets are modelled as serial numbers; they must all be increasing numbers not exceeding 1000. \mathcal{G}_T below specifies this scenario:

$$\begin{aligned} \mathcal{G}_T &= \text{C} \rightarrow \text{S} : (x : \text{int}) \{x \geq 0 \wedge \text{C.credit} \geq x\} \langle \text{C.credit} := \text{C.credit} - x \rangle. \\ &\quad \text{S} \rightarrow \text{C} : \{ \text{ok}(y : \text{int}) \{ \text{S.count} < 1000 \wedge y = \text{S.count} \} \langle E_{\text{ok}} \rangle . \text{end}, \\ &\quad \quad \text{ko} \langle \text{C.credit} := \text{C.credit} + x \rangle . \text{end} \} \\ E_{\text{ok}} &= \text{S.credit} := \text{S.credit} + x, \text{S.count} := \text{S.count} + 1 \end{aligned}$$

C has state variable `credit`, and S has state variables `credit` and `count` (a counter for serial numbers). The first interaction requires that the offer x does not exceed C's credit, and decrements the credit by x . S selects one of the two branches by either label `ok` or `ko`. The former branch can be selected only if `S.count < 1000`.

We denote with $\text{var}(\mathcal{G})$ the set of (interaction/state) variables and recursion parameters in \mathcal{G} , and with $\text{var}(A)$ the free variables of A (same for e). The set of variables that $p \in \mathcal{G}$ knows, written $\text{var}(\mathcal{G}) \upharpoonright p$, consists of: (i) the state variables of the form $p.x$ for some x , (ii) the interaction variables sent or received by p in \mathcal{G} , and (iii) the parameters of the recursive definitions $\mu t \langle y : A' \rangle (x : S) \{ A \} . \mathcal{G}'$ in \mathcal{G} such that p knows all the free variables in initialisation A' , and all free variables in A'' for all $t \langle y : A'' \rangle$ in \mathcal{G}' (we assume each recursion parameter known by exactly two participants).

Well-assertedness Our theory relies on two consistency principles: *history-sensitivity* and *temporal-satisfiability*. These principles were first introduced in [5]; we discuss them here as their adaptation to our stateful scenario requires non-trivial extensions.

By history-sensitivity each role must have enough information to fulfil the specified obligations, namely it requires that: (1) each role p knows all free variables in the predicates that p must guarantee, and (2) each role has enough information to perform the prescribed updates, that is (i) when to make an update, and (ii) which values to assign.

Definition 2 (History-sensitivity). \mathcal{G} is history-sensitive if for each interaction, of the form $p \rightarrow q : \{ l_i(x_i : U_i) \{ A_i \} \langle E_i \rangle . \mathcal{G}_i \}_{i \in I}$, occurring in \mathcal{G} , for all $i \in I$:

1. $\text{var}(\mathcal{G}) \upharpoonright p \supseteq \text{var}(A_i)$ (i.e., p knows all variables in $\text{var}(A_i)$),
2. for all $r.x := e$ in E_i : (i) either $r = p$ or $r = q$, and (ii) $\text{var}(\mathcal{G}) \upharpoonright r \supseteq \text{var}(e)$.

A checker for history-sensitivity can be found in the online report [4]. Consider the assertions:

$$\begin{aligned}\mathcal{G} &= p \rightarrow q : (x : \text{int}). q \rightarrow r : (y : \text{int}). r \rightarrow s : (z : \text{int})\{z > x\} \\ \mathcal{G}' &= p \rightarrow q : (y : \text{int}). q \rightarrow r : \{\text{OK}(w : \text{int})\langle r.x_1 := y, p.x_2 := y \rangle, \text{KO}\}\end{aligned}$$

\mathcal{G} violates (1) because r has to send a value for z that is greater than x without knowing x . \mathcal{G}' violates both clauses of (2): (i) because p must update x_2 not knowing whether and when the update should be done, and (ii) because in the second interaction r has to update x_1 with y without knowing y .³

By temporal-satisfiability, for each participant $p \in \mathcal{G}$, whenever it is p 's turn to send a value, p can find at least one selection branch and one value which satisfies the specified constraint. Temporal satisfiability is defined (and checked) using a function $\text{ts}(\mathcal{G}, A)$ which returns **true** only if \mathcal{G} always allows a path of interactions going through \mathcal{G} in *any possible state*. Considering all possible states makes the specification robust with respect to arbitrary interactions the same principal may be engaged in through other sessions. Predicate A is incrementally built as a conjunction of the predicates that appear in \mathcal{G} in all the recursive invocations and models the current set of assumptions.

Definition 3 (Temporal-satisfiability). Let \mathcal{G} be a global specification, and A a predicate. $\text{ts}(\mathcal{G}, A)$ is given by:

1. $\text{ts}(p \rightarrow q : \{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle. \mathcal{G}_i\}_{i \in I}, A) = \begin{cases} \bigwedge_{i \in I} \text{ts}(\mathcal{G}_i, A \wedge \underline{A}_i) & \text{if } A \supset \bigvee_{i \in I} \exists x_i. A_i \\ \text{false} & \text{otherwise} \end{cases}$
2. $\text{ts}(\mathcal{G}_1 \mid \mathcal{G}_2, A) = \text{ts}(\mathcal{G}_1, A) \wedge \text{ts}(\mathcal{G}_2, A)$
3. $\text{ts}(\mu t \langle e \rangle (x : S)\{A'\}. \mathcal{G}', A) = \begin{cases} \text{ts}(\mathcal{G}', A \wedge A') & \text{if } A \supset (A'[e/x]) \\ \text{false} & \text{otherwise} \end{cases}$
4. $\text{ts}(t_{A'(x)} \langle e \rangle, A) = \begin{cases} \text{true} & \text{if } A \supset A'[e/x] \\ \text{false} & \text{otherwise} \end{cases}$
5. $\text{ts}(\text{end}, A) = \text{true}$

\mathcal{G} satisfies temporal satisfiability if $\text{ts}(\mathcal{G}, \text{true}) = \text{true}$.⁴

In (1) the first condition for “if” demands that there exists at least one branch for which it is possible to find a value for x_i that satisfies the current predicate A_i . The function is called recursively extending the set of preconditions A with with the closure \underline{A}_i of predicate A_i (see Remark 4 below). (2) demands both parts of the composition are satisfiable. (3) and (4) check recursion, the latter relying on the annotation of recursive calls with the invariants of the corresponding recursive definitions.

³ [6] proposes algorithms to amend assertions that violate history-sensitivity and temporal-satisfiability as in [5]. No such algorithms have yet been investigated for the definitions introduced in this paper. Although relevant, the issue of amending inconsistent assertions is out of the scope of the current work.

⁴ This property can be relaxed by starting from a stronger precondition A as long as A is then implied by the principal invariants (which are defined in § 4).

Remark 4. The closure of a predicate A in \mathcal{G} , written \underline{A} , is the predicate obtained by closing with existential quantifiers the free state variables of \mathcal{G} in A . Whereas the values of interaction variables in a session do not change after they are introduced⁵, state variables can be updated a number of times in a number of different ways. Hence a predicate on state variables may be true at a certain time, and become false at a later time. Hereafter we use \underline{A} when we want to ‘keep’ only the persistent assumptions (those on interaction variables) of A .

The following global specification violates temporal satisfiability

$$p \rightarrow q : (x : \text{int})\{x > 0\}.q \rightarrow p : (y : \text{int})\{y = x \wedge y > 100\}$$

In fact, in the first interaction p is allowed to choose any positive value for x , for instance 10. In this case, q cannot find any value for y such that $y = 10 \wedge y > 100$.

Proposition 5. *Given a global assertion \mathcal{G} , let m be the size of the syntactic tree of \mathcal{G} , n be the maximum number of variables occurring in each predicate in \mathcal{G} , and $\text{eval}(A)$ be the complexity of predicate evaluation (if decidable). History-sensitivity can be checked in $O(m \times n)$. Temporal-satisfiability is decidable if predicate evaluation is decidable and, if decidable, it can be checked in $O(m) \times \text{eval}(A)$.*

Hereafter, we assume assertions to be *well-asserted*.

Local Assertions Each local assertion \mathcal{L} refers to a specific role. Syntax is given in the right part of Figure 1. Selection **(L-sel)** $p!\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle.\mathcal{L}_i\}_{i \in I}$ models an interaction where the role sends p a branch label l_i and a message x_i . A_i and E_i are the predicate and update respectively. **(L-bra)** is the dual branching. The others are as in the global assertions, except that a local assertion cannot be multi-threaded.

Given a global assertion \mathcal{G} , we can automatically derive the local assertions for each role $p \in \mathcal{G}$ by *projection*. The projection rules rely on a few auxiliary definitions: projection of a predicate, and projection of an update. The *projection of a predicate* A on p in \mathcal{G} , written $A \upharpoonright p$, is defined as $\exists \tilde{x}. A$ where $\tilde{x} = \text{var}(A) \setminus (\text{var}(\mathcal{G}) \upharpoonright p)$ (i.e., the existential closure of the variables that p does not know). The *projection of an update* E on p in \mathcal{G} , written $E \upharpoonright p$ is the update E' containing only the assignments $p_j.x_i := e_j$ such that $p_j = p$.

The projection rules for global assertions are as in [5], except that updates are now considered; their detailed presentation is not necessary to understand the results in this paper, hence we only give an illustration through Example 6. Henceforth, in $\mathcal{G} \upharpoonright p$ we shall omit the $p.$ prefix when referring to p 's state variables.

⁵ Actually, interaction variables in a recursion body are reused at each iteration. However, their values are due to follow the same constraints at each iteration.

Example 6. \mathcal{L}_C (resp. \mathcal{L}_S) is the projection of \mathcal{G}_T from Example 1 on C (resp. S).

$$\begin{aligned}
\mathcal{L}_C &= \mathbf{S}!(x : \text{int})\{x \geq 0 \wedge \text{credit} \geq x\}\langle \text{credit} := \text{credit} - x \rangle.\mathcal{L}'_C \\
\mathcal{L}'_C &= \mathbf{S}?\{\text{ok}(y : \text{int})\{\exists \text{S.count.S.count} < 1000 \wedge y = \text{S.count}\}.\text{end}, \\
&\quad \text{ko}\langle \text{credit} := \text{credit} + x \rangle.\text{end}\} \\
\mathcal{L}_S &= \mathbf{C}?(x : \text{int})\{\exists \text{C.credit}.x \geq 0 \wedge \text{C.credit} \geq x\}.\mathcal{L}'_S \\
\mathcal{L}'_S &= \mathbf{C}!\{\text{ok}(y : \text{int})\{\text{count} < 1000 \wedge y = \text{count}\} \\
&\quad \langle \text{credit} := \text{credit} + x, \text{count} := \text{count} + 1 \rangle.\text{end}, \\
&\quad \text{ko}.\text{end}\}
\end{aligned}$$

The projection of the first interaction of \mathcal{G}_T on sender C (resp. receiver S) is a send/select (resp. a receive/branch). The predicates/updates of the projections on a role are the projections of the predicates/updates on that role.⁶ The continuation is projected similarly, proceeding point-wise for each branch. Sometimes the projected predicate includes information about constraints of interactions between third parties (without however revealing the actual values exchanged by the third parties), e.g., $\exists \text{S.count.S.count} < 1000 \wedge y = \text{S.count}$ provides C with precondition $y < 1000$.

Well-assertedness is easily extended to local assertions.

3 Multiparty networks with local states

We consider networks of interactional entities called *principals* linked by a common global transport, modelled as queues. Each principal runs a *located process*, that is a process with multiparty session primitives [1, 18] (to enable rigorous representation of conversation structures) and with a *local state*.

Syntax The syntax of networks and processes is given in Figure 2 and is a refined version of the multiparty session π -calculus from [1, 10] with local states. A local state σ maps a signature $[\tilde{x} : \tilde{S}]$ of typed pairwise disjoint state variables \tilde{x} to their sorts. We use the injective function $\text{id}(\sigma)$ to map each local state to an identifier.

A network can be an empty network \emptyset , a located process $[P]\sigma$, a parallel composition of networks $N_1 \mid N_2$, a new session name $(\nu s)N$ which binds s in N , or a queue $s : h$ where h are messages in transit through session channel s . A network is *initial* if it has no new session names and queues, otherwise it is *runtime*. We denote the free session channels in N with $\text{fn}(N)$, similarly for P with $\text{fn}([P]\sigma) = \text{fn}(P)$.

Session request (**P-req**) multicasts a request to each session accept process (**P-acc**), e.g., $a[i](y).P$ with $i \in \{2, \dots, n\}$, by synchronisation through a shared name a and continuing as P . (**P-sel**) is Dijkstra's guarded command [15] and (**P-bra**) is the branching process; they represent communications through an established session k . (**P-sel**) acts as role p in session k and sends role q one of the labels l_i . The choice of the label is determined by boolean expressions e_i , assuming $\bigvee_{i \in I} e_i = \text{true}$ and $i \neq j$ implies

⁶ Note that by well-assertedness (clause 1) the projection of a predicate on the sender of an interaction is always the predicate itself.

(network)	$N ::= \emptyset \mid [P]\sigma \mid N_1 N_2 \mid (\nu s)N \mid s : h$		
(state/queue/value)	$\sigma ::= [\tilde{x} : \tilde{S}] \mapsto \tilde{S} \quad h ::= \emptyset \mid (p, q, l\langle v \rangle) \cdot h \quad v ::= n \mid s[p]$		
(process)	$P ::= \bar{a}[n](y).P$	(P-req)	$\mid P \mid Q$ (P-par)
	$\mid a[i](y).P$	(P-acc)	$\mid (\mu X(x).P)\langle e \rangle$ (P-def)
	$\mid k[p, q]!\{e_i \mapsto l_i\langle e'_i \rangle(x_i)\langle E_i \rangle; P_i\}_{i \in I}$	(P-sel)	$\mid X\langle e \rangle$ (P-call)
	$\mid k[p, q]?\{l_i(x_i)\langle E_i \rangle.P_i\}_{i \in I}$	(P-bra)	$\mid \mathbf{0}$ (P-idle)
(channel/update/exp)	$k ::= y \mid s \quad E ::= \emptyset \mid E; x := e \quad e ::= v \mid e \text{ op } e$		
x, y, \dots	interaction variables	x, y, \dots	state variables
a, b, \dots	shared name	X, Y, \dots	process variables
		n, n', \dots	constants

Fig. 2. Syntax of networks and processes

$e_i \wedge e_j = \text{false}$. Each label l_i is sent with the corresponding expression e'_i which specifies the value for x_i , assuming e'_i and x_i have the same type.⁷ **(P-bra)** plays role q in session k and is ready to receive from p one of the labels l_i and a value for the corresponding x_i , then behaves as P_i after instantiating x_i with the received value. In guarded command (resp. branching), the local state of the sender (resp. receiver) is updated according to update E_i ; in both processes each x_i binds its occurrences in P_i and E_i .

Example 7. Processes P_S and P_C implement \mathcal{L}_S and \mathcal{L}_C , respectively, from Example 6.

$$\begin{aligned}
P_S &= a[2](z).z[\mathbf{C}, \mathbf{S}]?(x); P'_S & E_{ok} &= \text{count} := \text{count} + 1, \text{credit} := \text{credit} + x \\
P'_S &= z[\mathbf{S}, \mathbf{C}]!\{\{\text{count} < 1000 \wedge x \geq 10\} \mapsto \text{ok}\langle \text{count} \rangle(y)\langle E_{ok} \rangle.\mathbf{0}, \\
&\quad \{\text{count} \geq 1000 \vee x < 10\} \mapsto \text{ko}.\mathbf{0}\} \\
P_C &= \bar{a}[2](w).w[\mathbf{C}, \mathbf{S}]!\langle 8 \rangle(x)\langle \text{credit} := \text{credit} - x \rangle; P'_C \\
P'_C &= w[\mathbf{S}, \mathbf{C}]?\{\text{ok}(y).\mathbf{0}, \text{ko}\langle \text{credit} := \text{credit} + x \rangle.\mathbf{0}\}
\end{aligned}$$

We let $\mathbf{C} = 1$ and $\mathbf{S} = 2$. P_S accepts a request to participate to a session specified by \mathcal{G}_T (assuming a has type \mathcal{G}_T) on channel z as role 2. In the established session z , the principal receives an offer x from the co-party. It follows a guarded command with two cases; if count has not reached its maximum value for serial numbers and the offer is greater than 10 then the first branch (**ok**) is taken and count is sent as y , otherwise the second branch (**ko**) is taken. Dually, P_C sends a request to participate to one instance of session \mathcal{G}_T as the role 1. A principal may repeatedly execute a process using recursion, or run concurrent instances of the same type of session (e.g., $[P_S \mid P_S]\sigma$) or different types of session (e.g., $[P_S \mid P_C]\sigma$) as discussed in Example 9.

Operational semantics The LTS is generated from the rules in Figure 3 using the following labels: $\ell ::= \bar{a}[n]\langle s \rangle \mid a[i]\langle s \rangle \mid s[p, q]!l\langle v \rangle \mid s[p, q]?l\langle v \rangle \mid \tau$. We denote with σ after E the state σ after the update E . We write $\sigma \models e \downarrow v$ for a closed expression e when it evaluates to v in σ .

⁷ Guarded command can be implemented using selection, if-then-else and lock-unlock. Although our theory is applicable to these primitives, we choose to make these low-level steps atomic for minimising the syntax.

$$\begin{array}{c}
\frac{[\bar{a}[n](y).P]\sigma \xrightarrow{\bar{a}[n]\langle s \rangle} [P[s/y]]\sigma \quad [a[i](y).P]\sigma \xrightarrow{a[i]\langle s \rangle} [P[s/y]]\sigma \quad (s \notin \text{fn}(P))}{[s[\mathbf{p}, \mathbf{q}]! \{e_i \mapsto l_i \langle e'_i \rangle (x_i) \langle E_i \rangle; P_i\}_{i \in I}]\sigma \xrightarrow{s[\mathbf{p}, \mathbf{q}]! l_j \langle v \rangle} [P[v/x_j]]\sigma'} \\
(j \in I \quad \sigma \models e'_j \downarrow v \quad \sigma \models e_j \quad \sigma' = \sigma \text{ after } E_j[v/x_j]) \\
[s[\mathbf{p}, \mathbf{q}]? \{l_i(x_i) \langle E_i \rangle, P_i\}_{i \in I}]\sigma \xrightarrow{s[\mathbf{p}, \mathbf{q}]? l_j \langle v \rangle} [P_j[v/x_j]]\sigma' \quad (j \in I \quad \sigma' = \sigma \text{ after } E_j[v/x_j]) \\
\frac{[P_1]\sigma_1 \xrightarrow{\bar{a}[n]\langle s \rangle} [P'_1]\sigma_1 \quad [P_i]\sigma_i \xrightarrow{a[i]\langle s \rangle} [P'_i]\sigma_i \quad (2 \leq i \leq n)}{[P_1]\sigma_1 \mid \dots \mid [P_n]\sigma_n \xrightarrow{\tau} (\nu s)(s : \emptyset \mid [P'_1]\sigma_1 \mid \dots \mid [P'_n]\sigma_n)} \\
\frac{[P]\sigma \xrightarrow{s[\mathbf{p}, \mathbf{q}]! l_j \langle v \rangle} [P']\sigma'}{[P]\sigma \mid s : h \xrightarrow{\tau} [P']\sigma' \mid s : h \cdot (\mathbf{p}, \mathbf{q}, l_j \langle v \rangle)} \quad \frac{[P]\sigma \xrightarrow{s[\mathbf{p}, \mathbf{q}]? l_j \langle v \rangle} [P']\sigma'}{[P]\sigma \mid s : (\mathbf{p}, \mathbf{q}, l_j \langle v \rangle) \cdot h \xrightarrow{\tau} [P']\sigma' \mid s : h}
\end{array}$$

Fig. 3. Labelled transition for networks

The first and second rule are for requesting and accepting a session initialisation. The guarded command checks if condition e_j is satisfied in the current state σ , and sends a message consisting of one of the labels l_j and an expression e'_j (which is evaluated to a value v in state σ), updates σ according to E_j , and behaves as $P[v/x_j]$. Branching is symmetric. The synchronous session initialisation creates a new queue. We omit the standard context/structural equivalence rules.

4 Proof system for multiparty session logic with virtual states

In this section we outline how to obtain the syntactic validation of networks, written $\Gamma \vdash N \triangleright \Sigma$, assuming processes typable, following [5]. The proof rules rely on the following environments:

$$\begin{array}{l}
\Gamma ::= \emptyset \mid \Gamma, a : \mathcal{G} \mid \Gamma, X : (x : S)\mathcal{L}_1 @ \mathbf{p}_1, \dots, \mathcal{L}_n @ \mathbf{p}_n, \quad \Delta ::= \emptyset \mid \Delta, s[\mathbf{p}] : \mathcal{L}, \\
\Sigma ::= \emptyset \mid \Sigma, [\Delta]\sigma
\end{array}$$

Γ maps shared names to global assertions and process variables to their parameters. If $\Gamma \vdash a : \mathcal{G}$ then a session specified by \mathcal{G} can be initiated by processes (via session request or accept) using a . By the standard kinding rules, we check if the same free variable appears in different global types in Γ , then they have the same sort. The mapping of process variables is for the validation of recursive assertions. Δ maps session channels/roles to local assertions. If $\Delta \vdash s[\mathbf{p}] : \mathcal{L}$ then a session is active (i.e., it has been initialized) on channel s for role \mathbf{p} ; \mathcal{L} specifies the (part of the) session that has still to be executed. Σ is the specification of a network; each $[\Delta]\sigma$ is the specification of a located process with the respective virtual state.

We also use an *assertion environment* \mathcal{C} , which is incrementally built by conjunction of the predicates and boolean expressions (i.e., the conditions of a guarded commands) occurring in the processes being validated, and models their assumptions. Hereafter, given a predicate A and an update E , we define $A \text{ after } E$ to be the predicate obtained by substituting, for each assignment $x := e$ in E , each occurrence of x in A with e .

Modelling cross-session properties: the principal invariant Given a located process $[P]\sigma$ in a network, we want to allow the architect to model stable properties (i.e., invariant) over the variables in σ on across multiple sessions. We call these properties *principal invariant* of $[P]\sigma$, that is a predicate (following the syntax for A in Figure 1) over the state variables of σ . Hereafter we assume there exists a function $\mathcal{I}(\sigma)$ that given a local state σ returns the principal invariant for σ . Principal invariants depend from the application domain, and the architect should define them prior to the verification.

Example 8. Consider a located process $[P_C \mid P_S]\sigma_p$ with P_C and P_S from Example 7. Assume we want to require that the credit is always non-negative (i.e., the principal does not contracts debts) and that the counter does not exceed the maximum number of tickets which is 1000. We can enforce these constraints by setting the principal invariant $\mathcal{I}(\sigma_p)$ to be $\text{credit} \geq 0 \wedge 0 \leq \text{count} \leq 1000$.

Proof rules Figure 4 illustrates the proof rules for initial networks and processes.

$$\begin{array}{c}
\frac{\text{id}(\sigma_p) = \text{id}(\sigma_a) \quad \sigma_p, \sigma_a \models \mathcal{I}(\sigma_p) \quad \mathcal{I}(\sigma_p) \wedge \mathcal{C}; \Gamma \vdash P \triangleright \Delta}{\mathcal{C}; \Gamma \vdash [P]\sigma_p \triangleright [\Delta]\sigma_a} \quad \text{[N1]} \\
\frac{(\Gamma', \Delta', \sigma') \ni (\Gamma, \Delta, \sigma) \quad \mathcal{C} \supset \mathcal{C}' \quad \mathcal{C}'; \Gamma' \vdash N \triangleright [\Delta']\sigma'}{\mathcal{C}; \Gamma \vdash N \triangleright [\Delta]\sigma} \quad \text{[N2]} \\
\frac{-}{\mathcal{C}; \Gamma \vdash \emptyset \triangleright \emptyset} \quad \frac{\mathcal{C}; \Gamma \vdash N \triangleright \Sigma \quad \mathcal{C}; \Gamma \vdash N' \triangleright \Sigma'}{\mathcal{C}; \Gamma \vdash N \mid N' \triangleright \Sigma, \Sigma'} \quad \text{[N3/N4]} \\
\hline
\frac{\mathcal{C}; \Gamma, a : \mathcal{G} \vdash P \triangleright y[\mathbf{i}] : \mathcal{G} \uparrow \mathbf{i}, \Delta}{\mathcal{C}; \Gamma, a : \mathcal{G} \vdash a[\mathbf{i}](y).P \triangleright \Delta} \quad \text{[MACC]} \\
\frac{\forall i \in I, \quad \mathcal{C} \wedge A_i; \Gamma \vdash P_i \triangleright \Delta, k[\mathbf{q}] : \mathcal{L}_i \quad \mathcal{C} \wedge A_i \supset (\mathcal{C} \text{ after } E_i)}{\mathcal{C}; \Gamma \vdash k[\mathbf{p}, \mathbf{q}]! \{l_i(x_i) \langle E_i \rangle, P_i\}_{i \in I} \triangleright \Delta, k[\mathbf{q}] : \mathbf{p}! \{l_i(x_i : U_i) \{A_i\} \langle E_i \rangle, \mathcal{L}_i\}_{i \in I}} \quad \text{[BCH]} \\
\frac{\forall i \in I \exists j \in J, \quad l_i = l_j \quad \mathcal{C} \wedge e_i; \Gamma \vdash P_i[e'_i/x_i] \triangleright \Delta, k[\mathbf{p}] : \mathcal{L}_j[e'_i/x_j] \quad \mathcal{C} \wedge e_i \supset (A_j \wedge (E_i = E_j) \wedge (\mathcal{C} \text{ after } E_j)) [e'_i/x_i]}{\mathcal{C}; \Gamma \vdash k[\mathbf{p}, \mathbf{q}]! \{e_i \mapsto l_i \langle e'_i \rangle (x_i) \langle E_i \rangle; P_i\}_{i \in I} \triangleright \Delta, k[\mathbf{p}] : \mathbf{q}! \{l_j(x_j : U_j) \{A_j\} \langle E_j \rangle, \mathcal{L}_j\}_{j \in J}} \quad \text{[SEL]} \\
\frac{\mathcal{C}; \Gamma \vdash P_1 \triangleright \Delta_1 \quad \mathcal{C}; \Gamma \vdash P_2 \triangleright \Delta_2}{\mathcal{C}; \Gamma \vdash P_1 \mid P_2 \triangleright \Delta_1, \Delta_2} \quad \frac{\Delta \text{ end only}}{\mathcal{C}; \Gamma \vdash \mathbf{0} \triangleright \Delta} \quad \text{[PAR/END]} \\
\frac{\mathcal{L}_1[e/x], \dots, \mathcal{L}_n[e/x] \text{ well-asserted}}{\mathcal{C}; \Gamma, X : (x)\mathcal{L}_1 @ \mathbf{p}_1, \dots, \mathcal{L}_n @ \mathbf{p}_n \vdash X \langle e \rangle \triangleright s[\mathbf{p}_1] : \mathcal{L}_1[e/x], \dots, s[\mathbf{p}_n] : \mathcal{L}_n[e/x]} \quad \text{[VAR]} \\
\frac{\mathcal{C}; \Gamma, X : (x)\mathcal{L}_1 @ \mathbf{p}_1, \dots, \mathcal{L}_n @ \mathbf{p}_n \vdash P \triangleright s[\mathbf{p}_1] : \mathcal{L}_1, \dots, s[\mathbf{p}_n] : \mathcal{L}_n}{\mathcal{C}; \Gamma \vdash (\mu X(x).P) \langle e \rangle \triangleright s[\mathbf{p}_1] : \mathcal{L}_1[e/x], \dots, s[\mathbf{p}_n] : \mathcal{L}_n[e/x]} \quad \text{[REC]}
\end{array}$$

Fig. 4. Proof rules for networks (top) and proof rules for processes (bottom)

[N1] decomposes the validation of a network into the validations of each located process against its corresponding specification Δ . The correspondence between principal and specification is checked by the clause $\text{id}(\sigma_p) = \text{id}(\sigma_a)$. Furthermore, local and virtual states must satisfy the principal invariant $\mathcal{I}(\sigma_p)$. P is then validated in the assertion environment extended (i.e., in conjunction with) the principal invariant.

[N2] is the rule for refinement. This rule is useful to validate processes even if they do not match exactly a given assertion as long as they implement a behaviour that is ‘more refined’ than the one prescribed. Refinement is also necessary to prove completeness of these rules (Theorem 14). We use the following refinement relation between specifications: $(\Gamma', \Delta', \sigma') \ni (\Gamma, \Delta, \sigma)$ if $(\Gamma', \Delta', \sigma')$ specifies a more refined behaviour than (Γ, Δ, σ) , in that it poses more restrictions on the output actions and poses less restrictions on the input actions. [N2] allows to refine the assertion environment \mathcal{C} by considering, in the premise, a weaker set of assumptions \mathcal{C}' .

[N3] is for empty networks and [N4] is for decomposing the validation of networks.

[M_{ACC}] validates a session accept on a shared channel a as role i provided that a is in the domain of Γ , and that the continuation P is validated against the specification Δ extended with the new session $y[i]$. In the (now active) session $y[i]$, P must behave as $\Gamma(a)$ projected on role i . The rule for session request is similar hence omitted.

[B_{CH}] validates the branching process. Δ must include an active session $k[q]$ on session channel k for the receiver role q . In the premise, the continuation for each branch i is required to be still valid in the assertion environment extended with A_i . In the second clause of the premise, for each branch i the update E_i must not invalidate \mathcal{C} ; this ensures that the update does not invalidate the principal invariant. The invariant is not mentioned explicitly (to keep the proof rules concise), but it is implied by \mathcal{C} . In fact, \mathcal{C} is the conjunction of (1) the principal invariant (by [N1]), (2) possibly some interaction predicates (by [B_{CH}]), and (3) possibly some boolean expressions (by [S_{EL}]). Since predicates (2), (3) and A_i do not contain free state variables⁸, then E_i can only invalidate the principal invariant (1); on the other hand (2), (3) and A_i are necessary premises (i.e., $\mathcal{C} \wedge A_i$ before the implication) as they may constrain interaction variables used by E_i .

In [S_{EL}] each branch i of the process must correspond to a branch j of the specification ($l_i = l_j$). The continuation must be validated in assertion environment \mathcal{C} extended with the closure \underline{e}_i of the condition of the branch e_i . The closure of boolean expression e_i is defined as the closure for predicates (see Remark 4). The clause at the bottom of the premise requires that, under the assumption $\mathcal{C} \wedge \underline{e}_i$: (1) expression e'_i satisfies A_j , (2) assertion and process have the same effects/updates on the states, (3) update E_j does not invalidate the principal invariant.⁹

[P_{AR}] is similar to [N2] but for parallel processes. [E_{ND}] validates the idle process provided that each active session in the specification Δ is of the form $y[p] : \text{end}$.

[V_{AR}] validates recursive call given that the active sessions in Δ correspond to the roles and local assertions associated to process variable X in Γ and that each \mathcal{L}_i is still well-asserted when the recursion parameter is substituted with e . [R_{EC}] is the standard rule for recursion definition. The validation of recursive processes is handled in a similar way to [5]; it uses a refinement rule for processes, similar to [N2] and omitted for simplicity, and the fact that assertions are refined by their unfolding. See [4] for more details.

⁸ By history-sensitivity A_i does not include any free state variable.

⁹ [B_{CH}]/[S_{EL}] can be extended to delegation adding the following clause for $U_i = \langle \mathcal{L} \rangle$: ([B_{CH}]) $\mathcal{C} \wedge A_i; \Gamma \vdash P_i \triangleright \Delta, k[q] : \mathcal{L}_i, x_i : \mathcal{L}$, and ([S_{EL}]) $\mathcal{C} \wedge \underline{e}_i; \Gamma \vdash P[e'_i/x_i] \triangleright \Delta', k[p] : \mathcal{L}_j[e'_j/x_j]$ with $\Delta = \Delta'', e_i : \mathcal{L}'_i$ and $\Delta' = \Delta''$.

Example 9. Consider the located process $[P_S \mid P_C]\sigma_p$ from Example 8 that executes two parallel threads: one selling a ticket and the other one buying another kind of ticket from another principal (the other principal is not modelled here). We show the validation of $\text{true}; \Gamma \vdash [P_S \mid P_C]\sigma_p \triangleright [\emptyset]\sigma_a$ proceeding top-down using the rules in Figure 4.

The global specification $[\emptyset]\sigma_a$ is initially empty since there are no active sessions. The active sessions will be added upon session request/accept by P_S and P_C . We assume $\sigma_p = \sigma_a = \{\text{count} : \text{int}, \text{credit} : \text{int}\} \mapsto \{10, 500\}$ and initially $\mathcal{C} = \text{true}$.

We first apply $[\text{N1}]$ with $\mathcal{I}(\sigma_p) = \text{credit} \geq 0 \wedge 0 \leq \text{count} \leq 1000$ from Example 8. For readability we will write \mathcal{I} instead of $\mathcal{I}(\sigma_p)$ in this example. Note that \mathcal{I} is satisfied by the local and virtual state. Next we apply rule $[\text{PAR}]$ that decomposes the derivation of two threads for P_S and P_C . We omit the illustration of the latter thread.

Below we illustrate the application of rule $[\text{MACC}]$ and $[\text{BCH}]$ to the former thread:

$$\frac{\frac{\mathcal{I} \wedge \{\exists \text{C.credit.C.credit} \geq x\}; \Gamma \vdash P'_S \triangleright z[\text{S}] : \mathcal{L}'_S \quad \mathcal{I} \wedge \{\exists \text{C.credit.C.credit} \geq x\} \supset (\mathcal{I} \text{ after } \emptyset)}{\mathcal{I}; \Gamma \vdash z[\text{C}, \text{S}]?(x).P'_S \triangleright z[\text{S}] : \mathcal{C}?(x : \text{Nat})\{\exists \text{C.credit.C.credit} \geq x\}.\mathcal{L}'_S} [\text{BCH}]}{\mathcal{I}; \Gamma \vdash P_S \triangleright \emptyset} [\text{MACC}]$$

For readability we will simplify $\mathcal{I} \wedge \{\exists \text{C.credit.C.credit} \geq x\}$ with the equivalent predicate $\bar{\mathcal{I}}$. Next, by $[\text{SEL}]$, setting $e = \text{count} < 1000 \wedge x \geq 10$, and $E_{ok} = \text{count} := \text{count} + 1, \text{credit} := \text{credit} + x$:

$$\frac{\frac{\mathcal{I} \wedge e \supset (\text{count} < 1000 \wedge y = \text{count} \wedge E_{ok} = E_{ok} \wedge \mathcal{I} \text{ after } E_{ok})[\text{count}/y] \quad \mathcal{I}; \Gamma \vdash \mathbf{0} \triangleright z[\text{S}] : \text{end} \quad \mathcal{I} \wedge \neg e \supset \text{true} \quad \mathcal{I}; \Gamma \vdash \mathbf{0} \triangleright z[\text{S}] : \text{end}}{\mathcal{I}; \Gamma \vdash z[\text{S}, \text{C}]!\{e \mapsto \text{ok}\langle \text{count} \rangle(y)\langle E_{ok} \rangle.\mathbf{0}, \neg e \mapsto \text{ko}.\mathbf{0}\} \vdash z[\text{S}] : \mathcal{C}!\{\text{ok}(y : \text{Nat})\{\text{count} < 1000 \wedge y = \text{count}\}\langle E_{ok} \rangle.\text{end}, \text{ko}.\text{end}\}}{\mathcal{I}; \Gamma \vdash z[\text{S}, \text{C}]!\{e \mapsto \text{ok}\langle \text{count} \rangle(y)\langle E_{ok} \rangle.\mathbf{0}, \neg e \mapsto \text{ko}.\mathbf{0}\} \vdash z[\text{S}] : \mathcal{C}!\{\text{ok}(y : \text{Nat})\{\text{count} < 1000 \wedge y = \text{count}\}\langle E_{ok} \rangle.\text{end}, \text{ko}.\text{end}\}} [\text{MACC}]$$

where each line in the premise refers to a branch (i.e., ok and ko). The most delicate clause is $\mathcal{I} \wedge e \supset (\text{count} < 1000 \wedge y = \text{count} \wedge E_{ok} = E_{ok} \wedge \mathcal{I} \text{ after } E_{ok})[\text{count}/y]$ which requires: (1) the interaction predicate to be satisfied under the current assumptions, and in fact $(\text{count} < 1000 \wedge y = \text{count})[\text{count}/y]$ is implied by e , (2) the updates to be consistent, and in fact trivially $E_{ok} = E_{ok}$, and (3) the update to not invalidate the invariant, and in fact $\text{credit} + x \geq 0 \wedge 0 \leq \text{count} + 1 \leq 1000$ is true under the assumptions $\text{credit} \geq 0, x \geq 10$ and $0 \leq \text{count}$. Finally we apply $[\text{END}]$ to the second premise of each branch.

The effectiveness of the proof rules depends on the logic chosen for the predicates, which depends on the application scenario. An example which fits these criteria is the Presburger arithmetic, which is often sufficiently expressive: practical uses of multiplication are encodable [17], and formulae with quantifiers may be calculated efficiently [20, 22].

Proposition 10. *The proof of $\mathcal{C}; \Gamma \vdash N \triangleright \Sigma$ is decidable if predicate evaluation is decidable.*

5 Soundness and completeness of the validation rules

We define a labelled transition relation for specifications $\langle \Gamma, \Sigma \rangle$ using the same labels as for networks. The main difference with the rules for networks is that predicates must

be satisfied for the transition to occur. We illustrate below the most remarkable rules (see [4] for the other rules). The rule for session request:

$$\langle (a : \mathcal{G}, \Gamma); [\Delta] \sigma \rangle \xrightarrow{\bar{a}[\mathbf{n}]\langle s \rangle} \langle (a : \mathcal{G}, \Gamma); [\Delta, s[1] : \mathcal{G} \uparrow 1] \sigma \rangle$$

extends Δ with the new session, given that $a : \mathcal{G}$ in Γ and the current state satisfies assertion invariant A . The rule for session accept is dual. The rule for selection/send:

$$\frac{j \in I \quad \sigma \models A_j[\mathbf{n}/x_j] \quad \sigma' = \sigma \text{ after } E_j[\mathbf{n}/x_j]}{\langle \Gamma; [\Delta, s[\mathbf{p}] : \mathbf{q}!\{l_i(x_i : U_i)\{A_i\}\langle E_i \rangle \mathcal{L}_i\}_{i \in I}] \sigma \rangle \xrightarrow{s[\mathbf{p}, \mathbf{q}]\langle \mathbf{n} \rangle} \langle \Gamma; [\Delta, s[\mathbf{p}] : \mathcal{L}_j[\mathbf{n}/x_j]] \sigma' \rangle}$$

moves to the continuation \mathcal{L}_j of the selected branch with the updated state σ' , given that the sent value \mathbf{n} satisfies predicate A_j for branch j in the current state σ .

Semantic conformance is defined using *conditional simulation* [5] to relate networks N to specifications $\langle \Gamma; \Sigma \rangle$.

Definition 11 (Conditional Simulation). A binary relation \mathcal{R} over N and $\langle \Gamma; \Sigma \rangle$ is a *conditional simulation* if, for each $(N, \langle \Gamma; \Sigma \rangle) \in \mathcal{R}$, if $N \xrightarrow{\ell} N'$ with ℓ being:

- (1) a branching then $\langle \Gamma; \Sigma \rangle$ is capable to move at the subject of ℓ , and if $\langle \Gamma; \Sigma \rangle \xrightarrow{\ell} \langle \Gamma; \Sigma' \rangle$ then $(N', \langle \Gamma; \Sigma' \rangle) \in \mathcal{R}$;
 - (2) a select, session request/accept, τ then $\langle \Gamma; \Sigma \rangle \xrightarrow{\ell} \langle \Gamma; \Sigma' \rangle$ and $(N', \langle \Gamma; \Sigma' \rangle) \in \mathcal{R}$.
- We write $N \lesssim \langle \Gamma; \Sigma \rangle$ if there exists a conditional simulation \mathcal{R} s.t. $(N, \langle \Gamma; \Sigma \rangle) \in \mathcal{R}$.

Conditional simulation is like standard simulation for all types of actions except for branching, for which it requires N to be simulated only for legal values/labels (i.e., a process must conform to a given specification as long as its environment does so).

Definition 12 (Satisfaction). N satisfies Σ in Γ and \mathcal{C} , written $\mathcal{C}; \Gamma \models N \triangleright \Sigma$, if for all closing substitutions $\tilde{\sigma}$ over N and Σ respecting Γ and \mathcal{C} , $N\tilde{\sigma} \lesssim \langle \Gamma; \Sigma\tilde{\sigma} \rangle$.

We write $\Gamma \models N \triangleright \Sigma$ when \mathcal{C} is `true` (e.g., for initial networks). Soundness and completeness for initial networks are stated below.

Theorem 13 (Soundness of Proof Rules). *Let N be an initial network. Then $\Gamma \vdash N \triangleright \Sigma$ implies $\Gamma \models N \triangleright \Sigma$.*

Theorem 14 (Completeness of Proof Rules). *Let $N \equiv \prod_{i \in I} [P_i] \sigma_{pi}$ be an initial network and $\Sigma = \prod_{i \in I} [\Delta_i] \sigma_{ai}$ be a specification. Assume that for all $i \in I$: (1) $\text{id}(\sigma_{pi}) = \text{id}(\sigma_{ai})$, (2) $\text{dom}(\sigma_{pi}) = \text{dom}(\sigma_{ai})$, and (3) $\mathcal{I}(\sigma_{pi})$ equivalent to `true`. If $\Gamma \models N \triangleright \Sigma$ then $\Gamma \vdash N \triangleright \Sigma$.*

(1-2) are for symmetry between N and Σ . (3) is necessary since the principals in N can make updates that differ from those made by the corresponding specifications in Σ ; this may not compromise the observable behaviour of N with respect to Σ , but N may invalidate some principal invariant which would make the thesis false.

6 Related work and further topics

The preceding integrations of session types with logical constraints include [13], based on concurrent constraints ensuring bi-linear usage of channels, and [5], based on logical

annotations on interactions, do not treat stateful properties. The combination of types and logical assertions referring to local state newly proposed in this paper enable fine-grained specifications and validation, which are not possible in [5, 13].

The expressiveness of the session type-based analyses has been greatly extended these past few years. On one side, the conversation calculus [8], contracts [11] and dynamic multirole session types [14] have opened the way to the modelling of protocols complex in their shapes, by describing more accurately how sessions can be joined or left, who is allowed participate. On the other side, works such as [5, 9] improved the way interactions inside a session are described: in [5], an assertion framework ensures logical properties on the communicated values, in [9], a security analysis guarantees that the coherence of the information flow is preserved. Our work improves the session type analyses in both directions: by proposing a division of the process being tested into separate principals that can join one or several sessions independently when conditions are matched and manage their own state, and by giving a description, inside each session, of the internal state of each participant and the property it should satisfy. A recent work [12] examines conditions to ensure that a stateful specification is robust w.r.t. asynchronous communications. Our work provides a complete proof system ensuring soundness for processes, whereas [12] only addresses properties of types.

The refinement types for channels (e.g. [3]) specify value dependency with logical constraints. For example, one might write $?(x : \text{int}, !\{y : \text{int} \mid y > x\})$ (using the notation from [16]). It specifies a dependency at a *single point* (channel). Our theory, based on multiparty sessions, can verify processes against a contract globally agreed by multiple distributed peers. [2] uses refinement types for channels to verify authentication in multiparty session protocols, but does not consider multi-session properties.

The work [7] investigates a relationship between a dual intuitionistic linear logic and binary session types, and shows that the former defines a proof system for a session calculus which can automatically characterise and guarantee a session fidelity and global progress. None of the above works treat either virtual states or logical specifications for interleaved multiparty sessions.

The use of Rely-Guarantee conditions or other related methods [19] instead of a single invariant does not increase the expressiveness of our system, but could ease proofs for parallel composition.

A future direction is to link between our static analysis and a dynamic monitor-based approach. Using our local specification as a monitor at each end-point, incoming and outgoing messages can be verified and filtered. We are currently working on this topic with [21] based on the logic developed in this paper.

References

1. L. Bettini et al. Global progress in dynamically interleaved multiparty sessions. In *CONCUR*, volume 5201 of *LNCS*, pages 418–433. Springer, 2008.
2. K. Bhargavan, R. Corin, P.-M. Deniérou, C. Fournet, and J. Leifer. Cryptographic protocol synthesis and verification for multiparty sessions. In *CSF*, pages 124–140, 2009.
3. K. Bhargavan, C. Fournet, and A. D. Gordon. Modular verification of security protocol code by typing. In *POPL*, pages 445–456, 2010.

4. L. Bocchi, R. Demangeon, and N. Yoshida. A multiparty multi-session logic (extended report). <http://www.cs.le.ac.uk/people/lb148/statefulassertions.html>.
5. L. Bocchi, K. Honda, E. Tuosto, and N. Yoshida. A theory of design-by-contract for distributed multiparty interactions. In *CONCUR*, volume 6269 of *LNCS*, pages 162–176, 2010.
6. L. Bocchi, J. Lange, and E. Tuosto. Three algorithms and a methodology for amending contracts for choreographies. *Scientific Annals of Computer Science*, 22(1):61–104, 2012.
7. L. Caires and F. Pfenning. Session types as intuitionistic linear propositions. In *CONCUR'10*, volume 6269 of *LNCS*, pages 222–236. Springer-Verlag, 2010.
8. L. Caires and H. T. Vieira. Conversation types. In *ESOP*, volume 5502 of *LNCS*, pages 285–300. Springer, 2009.
9. S. Capecchi, I. Castellani, and M. Dezani-Ciancaglini. Information flow safety in multiparty sessions. In *EXPRESS*, volume 64 of *EPTCS*, pages 16–30, 2011.
10. M. Carbone, K. Honda, and N. Yoshida. Structured interactional exceptions in session types. In *CONCUR*, volume 5201 of *LNCS*, pages 402–417. Springer, 2008.
11. G. Castagna and L. Padovani. Contracts for mobile processes. In *CONCUR 2009*, number 5710 in *LNCS*, pages 211–228, 2009.
12. T.-C. Chen and K. Honda. Specifying stateful asynchronous properties for distributed programs. (to appear in *CONCUR*), 2012.
13. M. Coppo and M. Dezani-Ciancaglini. Structured communications with concurrent constraints. In *TGC*, pages 104–125, 2008.
14. P.-M. Deniérou and N. Yoshida. Dynamic multirole session types. In *POPL*, pages 435–446, 2011.
15. E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM*, 18:453–457, August 1975.
16. T. Freeman and F. Pfenning. Refinement types for ML. *SIGPLAN Not.*, 26(6):268–277, 1991.
17. M. K. Ganai. Efficient decision procedure for bounded integer non-linear operations. In *HVC '08*, pages 68–83. LNCS, 2009.
18. K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL'08*, pages 273–284. ACM, 2008.
19. C. Jones. Abstraction for concurrency. In *SEFM*, LNCS, 2012. to appear.
20. G. Nelson and D. C. Oppen. A simplifier based on efficient decision algorithms. In *POPL'78*, pages 141–150. ACM, 1978.
21. Ocean Observatories Initiative (OOI). <http://www.oceanleadership.org/programs-and-partnerships/ocean-observing/ooi/>.
22. W. Pugh. The omega test: a fast and practical integer programming algorithm for dependence analysis. In *Supercomputing '91*, pages 4–13, New York, NY, USA, 1991. ACM.