

Monitoring Networks through Multiparty Session Types [★]

Laura Bocchi¹, Tzu-Chun Chen², Romain Demangeon², Kohei Honda², and Nobuko Yoshida³

¹ University of Leicester

² Queen Mary, University of London

³ Imperial College London

Abstract. In large-scale distributed infrastructures, applications are realised through communications among distributed components. The need for methods for assuring safe interactions in such environments is recognized, however the existing frameworks, relying on centralised verification or restricted specification methods, have limited applicability. This paper proposes a new theory of *monitored* π -calculus with dynamic usage of *multiparty session types* (MPST), offering a rigorous foundation for safety assurance of distributed components which asynchronously communicate through multiparty sessions. Our theory establishes a framework for semantically precise decentralised run-time enforcement and provides reasoning principles over monitored distributed applications, which complement existing static analysis techniques. We introduce asynchrony through the means of explicit routers and global queues, and propose novel equivalences between networks, that capture the notion of interface equivalence, i.e. equating networks offering the same services to a user. We illustrate our static-dynamic analysis system with an ATM protocol as a running example and justify our theory with results: satisfaction equivalence, local/global safety and transparency, and session fidelity.

1 Introduction

One of the main engineering challenges for distributed systems is the comprehensive verification of distributed software without relying on ad-hoc and expensive testing techniques. Multiparty session types (MPST) is a typing discipline for communication programming, originally developed in the π -calculus [14, 1, 2, 10, 11, 6] towards tackling this challenge. The idea is that applications are built starting from units of design called sessions. Each type of session, involving multiple roles, is first modelled from a global perspective (*global type*) and then projected onto *local types*, one for each role involved. As a verification method, the existing MPST systems focus on static type checking of endpoint processes against local types. The standard properties enjoyed by well-typed processes are

[★] This work has been partially sponsored by the project Leverhulme Trust Award Tracing Networks, Ocean Observatories Initiative and EPSRC EP/K011715/1, EP/G015635/1 and EP/G015481/1.

communication safety (all processes conform to globally agreed communication protocols) and freedom from deadlocks.

The direct application of the theoretical MPST techniques to the current practice, however, presents a few obstacles. Firstly, the existing type systems are targeted at calculi with first class primitives for linear communication channels and communication-oriented control flow; the majority of mainstream engineering languages would need to be extended in this sense to be suitable for syntactic session type checking. Unfortunately, it is not always straightforward to add these features to the specific host languages (e.g. linear resource typing for a very liberal language like C). Furthermore, the executable processes in a distributed system may be implemented in different languages. Secondly, for domains where dynamically typed or untyped languages are popular (e.g., Web programming), or in multi-organizational scenarios, the introduction of static typing infrastructure to support MPST may not be realistic.

This paper proposes a theoretical system addressing the above issues by enabling both static *and* dynamic verification of communicating processes. The aim is to capture the decentralised nature of distributed application development, providing better support for heterogeneous distributed systems by allowing components to be independently implemented, using different languages, libraries and programming techniques, as well as being independently verified, either statically or dynamically, while retaining the strong global safety properties of statically verified homogeneous systems.

This work is motivated in part by our ongoing collaboration with the Ocean Observatories Initiative (OOI) [16], a project to establish cyberinfrastructure for the delivery, management and analysis of scientific data from a large network of ocean sensor systems. Their architecture relies on the combination of high-level protocol specifications (to express how the infrastructure services should be used) and distributed run-time monitoring to regulate the behaviour of third-party applications in the system.

A formal theory for static/dynamic verification Our framework is based on the idea that, if each endpoint is *independently* verified (statically or dynamically) to conform to their local protocols, then the global protocol is respected *as a whole*. To this goal, we propose a new formal model and bisimulation theories of heterogeneous networks of monitored and unmonitored processes.

For the first time, we make explicit the *routing mechanism* implicitly present inside the MPST framework: in a session, messages are sent to abstract roles (e.g. to a Seller) and the router, a dynamically updated component of the network, translates these roles into actual addresses.

By taking this feature into account when designing novel equivalences, our formal model can relate networks built in different ways (through different distributions or relocations of services) but offering the same *interface* to an external observer. The router, being in charge of associating roles with principals, hides to an external user the internal composition of a network: what distinguishes two networks is not their structure but the services they are able to perform, or more precisely, the local types they offer to the outside.

We formally define a satisfaction relation to express when the behaviour of a network conforms to a global specification and we prove a number of properties of our model. *Local safety* states that a monitored process respects its local protocol, i.e. that dynamic verification by monitoring is sound, while *local transparency* states that a monitored process has equivalent behaviour to an unmonitored but well-behaved process, e.g. statically verified against the same local protocol. *Global safety* states that a system satisfies the global protocol, provided that each participant behaves as if monitored, while *global transparency* states that a fully monitored network has equivalent behaviour to an unmonitored but well-behaved network, i.e. in which all local processes are well-behaved against the same local protocols. *Session fidelity* states that, as all message flows of a network satisfy global specifications, whenever the network changes because some local processes take actions, all message flows continue to satisfy global specifications. Together, these properties justify our framework for decentralised verification by allowing monitored and unmonitored processes to be safely mixed while preserving protocol conformance for the entire network. Technically, these properties also ensure the coherence of our theory, by relating the satisfaction relations with the semantics and static validation procedures.

Paper summary and contributions §2 introduces the formalisms for protocol specifications (§2.1) and networks (§2.2) used to provide a formal framework for monitored networks based on π -calculus processes and protocol-based runtime enforcement through monitors. §3 introduces: a semantics for specifications (§3.1), a novel behavioural theory for compositional reasoning over monitored networks through the use of equivalences (bisimilarity and barbed congruence) and the satisfaction relation (§3.2). §3.4 establishes key properties of monitored networks, namely local/global safety, transparency, and session fidelity. We discuss future and related work in §4. The proofs can be found in the Appendices.

2 Types, Processes and Networks: a Formal Presentation

This section and the next one provide a theoretical basis for protocol-centred safety assurance. We first summarise the syntax of MPSTs (multiparty session types) annotated with logical assertions [2]. We then introduce a novel *monitored session calculus* as a variant of the π -calculus, modelling distributed dynamic components (whose behaviours are realised by processes) and monitors, all residing in global networks.

2.1 Multiparty Session Types with Assertions

Multiparty session types with assertions [2] are abstract descriptions of the structure of interactions among the participants of a multiparty session, specifying potential flows of messages, the conditions under which these interactions may be done, and the constraints on the communicated values. In this framework, *global types with assertions*, or just *global types*, describe multiparty sessions

from a network perspective. From global types one can derive, through *endpoint projection*, *local types with assertions*, or just *local types*, describing the protocol from the perspective of a single endpoint.

$$\begin{aligned}
A &::= \text{tt} \mid \text{ff} \mid e_1 = e_2 \mid e_1 < e_2 \mid \neg A \mid A_1 \wedge A_2 \mid A_1 \vee A_2 \\
e &::= v \mid e_1 + e_2 \mid e_1 - e_2 \mid e_1 * e_2 \mid e_1 \bmod e_2 \quad S ::= \text{bool} \mid \text{int} \mid \text{string} \\
G &::= \mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I} \mid G_1 \mid G_2 \mid G_1 ; G_2 \mid \mu t. G \mid \mathbf{t} \mid \epsilon \mid \text{end} \\
T &::= \mathbf{r}!\{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I} \mid \mathbf{r}?\{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I} \mid T_1 \mid T_2 \mid T_1 ; T_2 \mid \\
&\quad \mu t. T \mid \mathbf{t} \mid \epsilon \mid \text{end}
\end{aligned}$$

The syntax of the global types (G, G', \dots) and local types (T, T', \dots) is given above. The grammar is based on [2, 11] extended with parallel threads, which also require sequential composition to merge parallel threads as in [18]. We let values v, v', \dots range over boolean constants, numerals and strings, and e, e', \dots range over first-order expressions. For expressing constraints, we use logical predicates, or *assertions*, ranged over by A, A', \dots , following the grammar given above, although other decidable logics could be used.¹ The *sorts* of exchanged values (S, S', \dots) consists of atomic types.

Global types with assertions $\mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I}$ models an interaction where role \mathbf{r}_1 sends role \mathbf{r}_2 one of the *branch labels* l_i , as well as a value denoted by an *interaction variable* x_i of sort S_i . Interaction variable x_i binds its occurrences in A_i and G_i . A_i is the assertion which needs to hold for \mathbf{r}_1 to select l_i , and which may constrain the values instantiating x_i . $G_1 \mid G_2$ specifies two parallel sessions, and $G_1 ; G_2$ denotes sequential composition (assuming that G_1 does not include *end*). $\mu t. G$ is a recursive type, where t is guarded in G in the standard way, ϵ is the inaction for absence of communication, and *end* ends the session.

Example 1 (ATM: the global type). We present global type G_{ATM} that specifies an *ATM* scenario. Each session of *ATM* involves three roles: a client (**C**), the payment server (**S**) and a separate authenticator (**A**).

$$\begin{aligned}
G_{\text{ATM}} &= \mathbf{C} \rightarrow \mathbf{A} : \{ \text{Login}(x_i : \text{string})\{\text{tt}\}. \\
&\quad \mathbf{A} \rightarrow \mathbf{S} : \{ \text{LoginOK}()\{\text{tt}\}. \mathbf{A} \rightarrow \mathbf{C} : \{ \text{LoginOK}()\{\text{tt}\}. G_{\text{Loop}}, \\
&\quad \quad \text{LoginFail}()\{\text{tt}\}. \mathbf{A} \rightarrow \mathbf{C} : \{ \text{LoginFail}()\{\text{tt}\}. \text{end} \} \} \} \\
G_{\text{Loop}} &= \mu \text{LOOP}. \\
&\quad \mathbf{S} \rightarrow \mathbf{C} : \{ \text{Account}(x_b : \text{int})\{x_b \geq 0\}. \\
&\quad \mathbf{C} \rightarrow \mathbf{S} : \{ \text{Withdraw}(x_p : \text{int})\{x_p > 0 \wedge x_b - x_p \geq 0\}. \text{LOOP}, \\
&\quad \quad \text{Deposit}(x_d : \text{int})\{x_d > 0\}. \text{LOOP}, \\
&\quad \quad \text{Quit}()\{\text{tt}\}. \text{end} \} \}
\end{aligned}$$

At the start of the session **C** sends its login details x_i to **A**, then **A** informs **S** and **C** whether the authentication is successful, by choosing either the branch

¹ We use a logic without quantifiers, contrary to [2], to simplify the presentation and because *monitorability*, defined later in this section, makes them unnecessary.

with label `LoginOK` or `LoginFail`. In the former case `C` and `S` enter a transaction loop specified by G_{Loop} . In each iteration `S` sends `C` the amount x_b available in the account, which must be non negative. Next, `C` has three choices: `Withdraw` withdraws an amount x_p (x_p must be positive and not exceed the current amount x_b) and repeats the loop, `Deposit` deposits a positive amount x_d in the account and repeats the loop, and `Quit` ends the session.

We consider global types that satisfy the consistency conditions defined in [10, 2, 11] which rule out, for instance, protocols where interactions have causal relations that cannot be enforced (e.g., we write $\mathbf{r}_A \rightarrow \mathbf{r}_B : \mathbf{l}_1() \{\text{tt}\} \mid \mathbf{r}_C \rightarrow \mathbf{r}_D : \mathbf{l}_2() \{\text{tt}\}$ instead of $\mathbf{r}_A \rightarrow \mathbf{r}_B : \mathbf{l}_1() \{\text{tt}\}. \mathbf{r}_C \rightarrow \mathbf{r}_D : \mathbf{l}_2() \{\text{tt}\}$). In addition we assume *monitorability* requiring that in all the interactions of the form $\mathbf{r} \rightarrow \mathbf{r}' : \mathbf{l}(x : S)\{A\}$ occurring in a global type G both \mathbf{r} and \mathbf{r}' *know* (i.e., have sent or received in a previous or in this interaction) the free variables in A .

Local types with assertions Each local type T is associated with a role taking part in a session. Local type $\mathbf{r}!\{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I}$ models an interaction where the role under consideration sends \mathbf{r} a branch label l_i and a message denoted by an interaction variable x_i of sort S_i . Its dual is the receive interaction $\mathbf{r}?\{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I}$. The other local types are similar to the global types.

One can derive a set of local types T_i from a global type G by *endpoint projection*, defined as in [2]. We write $G \upharpoonright \mathbf{r}$ for the projection of G onto role \mathbf{r} . We illustrate the main projection rule, which is for projecting a global type modelling an interaction. Let G be $(\mathbf{r} \rightarrow \mathbf{r}' : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I})$; the projection of G on \mathbf{r} is $\mathbf{r}!\{l_i(x_i : S_i)\{A_i\}.(G_i \upharpoonright \mathbf{r})\}_{i \in I}$, and the projection of G on \mathbf{r}' is $\mathbf{r}'?\{l_i(x_i : S_i)\{A_i\}.(G_i \upharpoonright \mathbf{r}')\}_{i \in I}$. The other rules are homomorphic, following the grammar of global types inductively.

Example 2 (ATM: the local type of C). We present the *local type* T_C obtained by projecting G_{ATM} on role `C`.

$$T_C = \mathbf{A}!\{\text{Login}(x_i : \text{string})\{\text{tt}\}. \mathbf{A}?\{\text{LoginOK}()\{\text{tt}\}. T_{\text{Loop}} \text{LoginFail}()\{\text{tt}\}. \text{end}\}\} \quad \left| \quad T_{\text{Loop}} = \mu \text{LOOP}. \mathbf{S}?\{\text{Account}(x_b : \text{int})\{x_b \geq 0\}. \mathbf{S}!\{\text{Withdraw}(x_p : \text{int})\{x_p > 0 \wedge x_b - x_p \geq 0\}. \text{LOOP}, \text{Deposit}(x_d : \text{int})\{x_d > 0\}.\text{LOOP}, \text{Quit}()\{\text{tt}\}.\text{end}\}\}$$

T_C specifies the behaviour that `C` should follow to meet the contract of global type G_{ATM} . T_C states that `C` should first authenticate with `A`, then receive the `Account` message from `S`, and then has the choice of sending `Withdraw` (and enact the recursion), or `Deposit` (and enact the recursion) or `Quit` (and end the session).

2.2 Formal Framework of Processes and Networks

In our formal framework, each distributed application consists of one or more sessions among *principals*. A principal with behaviour P and name α is represented as $[P]_\alpha$. A *network* is a set of principals together with a (unique) *global*

transport, which abstractly represents the communication functionality of a distributed system. The syntax of processes, principals and networks is given below, building on the multiparty session π -calculus from [1].

$$\begin{aligned}
P &::= \bar{a}\langle s[\mathbf{r}] : T \rangle \mid a(y[\mathbf{r}] : T).P \mid k[\mathbf{r}_1, \mathbf{r}_2]!l\langle e \rangle \mid k[\mathbf{r}_1, \mathbf{r}_2]? \{l_i(x_i).P_i\}_{i \in I} \mid \\
&\quad \text{if } e \text{ then } P \text{ else } Q \mid P \mid Q \mid \mathbf{0} \mid \mu X.P \mid X \mid P; Q \mid (\nu a)P \mid (\nu s)P \\
N &::= [P]_\alpha \mid N_1 \mid N_2 \mid \mathbf{0} \mid (\nu a)N \mid (\nu s)N \mid \langle r ; h \rangle \\
r &::= a \mapsto \alpha \mid s[\mathbf{r}] \mapsto \alpha \quad h ::= m \cdot h \mid \emptyset \quad m ::= \bar{a}\langle s[\mathbf{r}] : T \rangle \mid s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle
\end{aligned}$$

$\mathbf{r}, \mathbf{r}_1, \dots$ roles	s, s', \dots session names	X, Y, \dots process variables
a, b, \dots shared names	x, y, \dots variables	P, Q, \dots processes
α, β, \dots principal names	N, N', \dots networks	

Processes Processes are ranged over by P, P', \dots and communicate using two types of channel: *shared channels* (or shared names) used by processes for sending and receiving invitations to participate in sessions, and *session channels* (or session names) used for communication *within* established sessions. One may consider session names as e.g., URLs or service names.

The *session invitation* $\bar{a}\langle s[\mathbf{r}] : T \rangle$ invites, through a shared name a , another process to play \mathbf{r} in a session s . The *session accept* $a(y[\mathbf{r}] : T).P$ receives a session invitation and, after instantiating y with the received session name, behaves in its continuation P as specified by local type T for role \mathbf{r} . The *selection* $k[\mathbf{r}_1, \mathbf{r}_2]!l\langle e \rangle$ sends, through session channel k (of an established session), and as a sender \mathbf{r}_1 and to a receiver \mathbf{r}_2 , an expression e with label l . The *branching* $k[\mathbf{r}_1, \mathbf{r}_2]? \{l_i(x_i).P_i\}_{i \in I}$ is ready to receive one of the labels and a value, then behaves as P_i after instantiating x_i with the received value. We omit labels when I is a singleton. The *conditional*, *parallel* and *inaction* are standard. The *recursion* $\mu X.P$ defines X as P . Processes $(\nu a)P$ and $(\nu s)P$ hide shared names and session names, respectively.

Principals and network A *principal* $[P]_\alpha$, with its process P and name α , represents a unit of behaviour (hence verification) in a distributed system. A *network* N is a collection of principals with a unique global transport.

A *global transport* $\langle r ; h \rangle$ is a pair of a routing table which delivers messages to principals, and a global queue. Messages between two parties inside a single session are ordered (as in a TCP connection), otherwise unordered. More precisely, in $\langle r ; h \rangle$, h is a *global queue*, which is a sequence of *messages* $\bar{a}\langle s[\mathbf{r}] : T \rangle$ or $s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle$, ranged over by m . These m represent messages-in-transit, i.e. those messages which have been sent from some principals but have not yet been delivered. The *routing table* r is a finite map from session-roles and shared names to principals. If, for instance, $s[\mathbf{r}] \mapsto \alpha \in r$ then a message for \mathbf{r} in session s will be delivered to principal α .

Let n, n', \dots range over shared and session channels. A network N which satisfies the following conditions is *well-formed*: (1) N contains at most one global transport; (2) two principals in N never have the same principal name;

and (3) if $N \equiv (\nu \tilde{n})(\prod_i [P_i]_{\alpha_i} | \langle r ; h \rangle)$, each free shared or session name in P_i and h occurs in \tilde{n} (we use $\prod_i P_i$ to denote $P_1 | P_2 \cdots | P_n$).

Semantics The reduction relation for dynamic networks is generated from the rules below, which model the interactions of principals with the global queue.

$$\begin{array}{l}
[\bar{a}\langle s[\mathbf{r}] : T \rangle]_{\alpha} | \langle r ; h \rangle \longrightarrow [\mathbf{0}]_{\alpha} | \langle r ; h \cdot \bar{a}\langle s[\mathbf{r}] : T \rangle \rangle \quad [\text{REQ}] \\
[a(y[\mathbf{r}] : T).P]_{\alpha} | \langle r ; \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot h \rangle \longrightarrow [P[s/y]]_{\alpha} | \langle r \cdot s[\mathbf{r}] \mapsto \alpha ; h \rangle^{\dagger} \quad [\text{ACC}] \\
[s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle]_{\alpha} | \langle r ; h \rangle \longrightarrow [\mathbf{0}]_{\alpha} | \langle r ; h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle \rangle \rangle^{\dagger\dagger} \quad [\text{SEL}] \\
[s[\mathbf{r}_1, \mathbf{r}_2]?\{l_i(x_i).P_i\}_i]_{\alpha} | \langle r ; s\langle \mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle \rangle \cdot h \rangle \longrightarrow [P_j[v/x_j]]_{\alpha} | \langle r ; h \rangle^{\dagger\dagger\dagger} \quad [\text{BRA}] \\
[\text{if tt then } P \text{ else } Q]_{\alpha} \longrightarrow [P]_{\alpha} \quad [\text{if ff then } P \text{ else } Q]_{\alpha} \longrightarrow [Q]_{\alpha} \quad [\text{CND}] \\
\frac{[P]_{\alpha} | N \longrightarrow [P']_{\alpha} | N'}{[\mathcal{E}(P)]_{\alpha} | N \longrightarrow [\mathcal{E}(P')]_{\alpha} | N'} \quad \frac{e \longrightarrow e'}{[\mathcal{E}(e)]_{\alpha} \longrightarrow [\mathcal{E}(e')]_{\alpha}} \quad \frac{N \longrightarrow N'}{\mathcal{E}(N) \longrightarrow \mathcal{E}(N')} \quad [\text{CTX}] \\
\uparrow : r(a) = \alpha \quad \uparrow\uparrow : r(s[\mathbf{r}_2]) \neq \alpha \quad \uparrow\uparrow\uparrow : r(s[\mathbf{r}_2]) = \alpha
\end{array}$$

$$\mathcal{E} ::= () \mid \mathcal{E} | P \mid (\nu s)\mathcal{E} \mid (\nu a)\mathcal{E} \mid \mathcal{E}; P \mid \mathcal{E} | N \mid \text{if } \mathcal{E} \text{ then } P \text{ else } Q \mid s[\mathbf{r}_1, \mathbf{r}_2]!l\langle \mathcal{E} \rangle$$

Rule [REQ] places an invitation in the global queue. Dually, in [ACC], a process receives an invitation on a shared name from the global queue, assuming a message on a is to be routed to α . As a result, the routing table adds $s[\mathbf{r}] \mapsto \alpha$ in the entry for s . Rule [SEL] puts in the queue a message sent from \mathbf{r}_1 to \mathbf{r}_2 , which selects label l_j and carries v , if it is not going to be routed to α (i.e. sent to self). Dually, [BRA] gets a message with label l_j from the global queue, so that the j -th process P_j receives value v . The reduction is also defined modulo the structural congruence \equiv defined by the standard laws over processes/networks, the unfolding of recursion ($\mu X.P \equiv P[\mu X.P/X]$) and the associativity and commutativity and the rules of message permutation in the queue [14, 10]. The other rules are standard.

Example 3 (ATM: an implementation). We now illustrate the processes implementing the client role of the *ATM* protocol. We let P_C be the process implementing T_C (from Example 2) and communicating on session channel s .

$$\begin{array}{l}
P_C = s[\mathbf{C}, \mathbf{A}]!\text{Login}(\text{alice_pwd}123); \\
\quad s[\mathbf{A}, \mathbf{C}]?\{\text{LoginOK}(); \mu X.P'_C, \text{LoginFail}().\mathbf{0}\} \\
P'_C = s[\mathbf{S}, \mathbf{C}]?\text{Account}(x_b); P''_C
\end{array} \left| \begin{array}{l}
P''_C = \text{if } \text{getmore}() \wedge (x_b \geq 10) \\
\quad \text{then } s[\mathbf{C}, \mathbf{S}]!\text{Withdraw}(10); X \\
\quad \text{else } s[\mathbf{C}, \mathbf{S}]!\text{Quit}(); \mathbf{0}
\end{array} \right.$$

Note that P_C selects only two of the possible branches (i.e., **Withdraw** and **Quit**) and **Deposit** is never selected. One can think of P_C as an ATM machine that only allows to withdraw a number of £10 banknotes, until the amount exceeds the current balance. This ATM machine does not allow deposits. We assume $\text{getmore}()$ to be a local function to the principal running P_C that returns **tt** if more notes are required (**ff** otherwise). P_S below implements the server role:

$$\begin{array}{l}
P_S = s[\mathbf{A}, \mathbf{S}]?\{\text{LoginOK}(); \mu X.P'_S, \text{LoginFail}().\mathbf{0}\} \\
P'_S = s[\mathbf{S}, \mathbf{C}]!\text{Account}(\text{getBalance}()); P''_S
\end{array} \left| \begin{array}{l}
P''_S = s[\mathbf{C}, \mathbf{S}]?\{\text{Withdraw}(x_p).X, \\
\quad \text{Deposit}(x_d).X, \\
\quad \text{Quit}().\mathbf{0}\}
\end{array} \right.$$

where $getBalance()$ is a local function to the principal running P_S that synchronously returns the current balance of the client.

3 Theory of Dynamic Safety Assurance

In this section we formalise the specifications (based on local types) used to guard the runtime behaviour of the principals in a network. These specifications can be embedded into system monitors, each wrapping a principal to ensure that the ongoing communication conforms to the given specification. Then, we present a behavioural theory for monitored networks and its safety properties.

3.1 Semantics of Global Specifications

The specification of the (correct) behaviour of a principal consists of an *assertion environment* $\langle \Gamma; \Delta \rangle$, where Γ is the *shared environment* describing the behaviour on shared channels, and Δ is the *session environment* representing the behaviour on session channels (i.e., describing the sessions that the principal is currently participating in). The syntax of Γ and Δ is given by:

$$\Gamma ::= \emptyset \mid \Gamma, a : \mathbf{I}(T[\mathbf{r}]) \mid \Gamma, a : \mathbf{O}(T[\mathbf{r}]) \quad \Delta ::= \emptyset \mid \Delta, s[\mathbf{r}] : T$$

In Γ , the assignment $a : \mathbf{I}(T[\mathbf{r}])$ (resp. $a : \mathbf{O}(T[\mathbf{r}])$) states that the principal can, through a , receive (resp. send) invitations to play role \mathbf{r} in a session instance specified by T . In Δ , we write $s[\mathbf{r}] : T$ when the principal is playing role \mathbf{r} of session s specified by T . Networks are monitored with respect to *collections of specifications* (or just *specifications*) one for each principal in the network. A specification Σ, Σ', \dots is a finite map from principals to assertion environments:

$$\Sigma ::= \emptyset \mid \Sigma, \alpha : \langle \Gamma; \Delta \rangle$$

The semantics of Σ is defined using the following labels:

$$\ell ::= \bar{a}(s[\mathbf{r}] : T) \mid a(s[\mathbf{r}] : T) \mid s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle \mid s[\mathbf{r}_1, \mathbf{r}_2]?l\langle v \rangle \mid \tau$$

The first two labels are for *invitation* actions, the first is for requesting and the second is for accepting. Labels with $s[\mathbf{r}_1, \mathbf{r}_2]$ indicate *interaction* actions for sending (!) or receiving (?) messages within sessions. The labelled transition relation for specification is defined by the rules below.

$$\begin{array}{c} \alpha : \langle \Gamma, a : \mathbf{O}(T[\mathbf{r}]); \Delta \rangle \xrightarrow{\bar{a}(s[\mathbf{r}] : T)} \alpha : \langle \Gamma, a : \mathbf{I}(T[\mathbf{r}]); \Delta \rangle \quad [\text{REQ}] \\ \frac{s \notin \text{dom}(\Delta)}{\alpha : \langle \Gamma, a : \mathbf{I}(T[\mathbf{r}]); \Delta \rangle \xrightarrow{a(s[\mathbf{r}] : T)} \alpha : \langle \Gamma, a : \mathbf{O}(T[\mathbf{r}]); \Delta, s[\mathbf{r}] : T \rangle} \quad [\text{ACC}] \\ \frac{\Gamma \vdash v : S_j, A_j[v/x_j] \downarrow \mathbf{tt}, j \in I}{\alpha : \langle \Gamma; \Delta, s[\mathbf{r}_2] : \mathbf{r}_1? \{l_i(x_i : S_i)\{A_i\}.T'_i\}_{i \in I} \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]?l_j\langle v \rangle} \alpha : \langle \Gamma; \Delta, s[\mathbf{r}_2] : T'_j[v/x_j] \rangle} \quad [\text{BRA}] \\ \frac{\Gamma \vdash v : S_j, A_j[v/x_j] \downarrow \mathbf{tt}, j \in I}{\alpha : \langle \Gamma; \Delta, s[\mathbf{r}_1] : \mathbf{r}_2! \{l_i(x_i : S_i)\{A_i\}.T'_i\}_{i \in I} \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle} \alpha : \langle \Gamma; \Delta, s[\mathbf{r}_1] : T'_j[v/x_j] \rangle} \quad [\text{SEL}] \\ \frac{\alpha : \langle \Gamma_1; \Delta_1 \rangle \xrightarrow{\ell} \alpha : \langle \Gamma'_1; \Delta'_1 \rangle}{\alpha : \langle \Gamma_1; \Delta_1 | \Delta_2 \rangle \xrightarrow{\ell} \alpha : \langle \Gamma'_1; \Delta'_1 | \Delta_2 \rangle} \quad \Sigma \xrightarrow{\tau} \Sigma \quad \frac{\Sigma_1 \xrightarrow{\ell} \Sigma_2}{\Sigma_1, \Sigma_3 \xrightarrow{\ell} \Sigma_2, \Sigma_3} \quad [\text{SPL, TAU, PAR}] \end{array}$$

Rule [REQ] allows α to send an invitation on a properly typed shared channel a (i.e., given that the shared environment maps a to $T[\mathbf{r}]$). Rule [ACC] allows α to receive an invitation to be role \mathbf{r} in a new session s , on a properly typed shared channel a . Rule [BRA] allows α , participating to sessions s as \mathbf{r}_2 , to receive a message with label l_j from \mathbf{r}_1 , given that A_j is satisfied after replacing x_j with the received value v . After the application of this rule the specification is T_j . Rule [SEL] is the symmetric (output) counterpart of [BRA]. We use \downarrow to denote the evaluation of a logical assertion. [SPL] is the parallel composition of two session environments where $\Delta_1|\Delta_2$ composes two local types: $\Delta_1|\Delta_2 = \{s[\mathbf{r}] : (T_1 \mid T_2) \mid T_i = \Delta_i(s[\mathbf{r}]), s[\mathbf{r}] \in \text{dom}(\Delta_1) \cap \text{dom}(\Delta_2)\} \cup \text{dom}(\Delta_1)/\text{dom}(\Delta_2) \cup \text{dom}(\Delta_2)/\text{dom}(\Delta_1)$. [TAU] says that the specification should be invariant under reduction of principals. [PAR] says if Σ_1 and Σ_3 are composable, after Σ_1 becomes as Σ_2 , they are still composable.

3.2 Semantics of Dynamic Monitoring

The endpoint monitor M, M', \dots for principal α is a specification $\alpha : \langle \Gamma; \Delta \rangle$ used to *dynamically* ensure that the messages to and from α are legal with respect to Γ and Δ . A *monitored network* N is a network N with monitors, obtained by extending the syntax of networks as:

$$N ::= N \mid M \mid N \mid N \mid (\nu s)N \mid (\nu a)N$$

The reduction rules for monitored networks are given below and use, in the premises, the labelled transitions of monitors. The labelled transitions of a monitor are the labelled transitions of its corresponding specification (given in § 3.1).

$$\begin{array}{c}
\text{[REQ]} \frac{M \xrightarrow{\bar{a}\langle s[\mathbf{r}]:T \rangle} M'}{[\bar{a}\langle s[\mathbf{r}]:T \rangle]_\alpha \mid M \mid \langle r; h \rangle \longrightarrow [\mathbf{0}]_\alpha \mid M' \mid \langle r; h \cdot \bar{a}\langle s[\mathbf{r}]:T \rangle \rangle} \\
\text{[ACC]} \frac{M \xrightarrow{a\langle s[\mathbf{r}]:T \rangle} M' \quad r(a) = \alpha}{[a(y[\mathbf{r}]:T).P]_\alpha \mid M \mid \langle r; \bar{a}\langle s[\mathbf{r}]:T \rangle \cdot h \rangle \longrightarrow [P[s/y]]_\alpha \mid M' \mid \langle r \cdot s[\mathbf{r}] \mapsto \alpha; h \rangle} \\
\text{[BRA]} \frac{M \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle} M' \quad r(s[\mathbf{r}_2]) = \alpha}{[s[\mathbf{r}_1, \mathbf{r}_2]?\{l_i(x_i).P_i\}_i]_\alpha \mid M \mid \langle r; s\langle \mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle \rangle \cdot h \rangle \longrightarrow [P_j[v/x_j]]_\alpha \mid M' \mid \langle r; h \rangle} \\
\text{[SEL]} \frac{M \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle} M' \quad r(s[\mathbf{r}_2]) \neq \alpha}{[s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle]_\alpha \mid M \mid \langle r; h \rangle \longrightarrow [\mathbf{0}]_\alpha \mid M' \mid \langle r; h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle \rangle} \\
\text{[REQER]} \frac{M \xrightarrow{\bar{a}\langle s[\mathbf{r}]:T \rangle} M'}{[\bar{a}\langle s[\mathbf{r}]:T \rangle]_\alpha \mid M \mid \langle r; h \rangle \longrightarrow [\mathbf{0}]_\alpha \mid M \mid \langle r; h \rangle} \\
\text{[ACCER]} \frac{M \xrightarrow{a\langle s[\mathbf{r}]:T \rangle} M'}{[a(y[\mathbf{r}]:T).P]_\alpha \mid M \mid \langle r; \bar{a}\langle s[\mathbf{r}]:T \rangle \cdot h \rangle \longrightarrow [a(y[\mathbf{r}]:T).P]_\alpha \mid M \mid \langle r; h \rangle} \\
\text{[SELER]} \frac{M \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle} M'}{[s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle]_\alpha \mid M \mid \langle r; h \rangle \longrightarrow [\mathbf{0}]_\alpha \mid M \mid \langle r; h \rangle}
\end{array}$$

The first four rules model reductions that are allowed by the monitor (i.e., in the premise). Rule [REQ] inserts an invitation in the global queue. Rule [ACC] is symmetric and updates the router so that all messages for role \mathbf{r} in session

s will be routed to α . Similarly, $\llbracket \text{BRA} \rrbracket$ (resp. $\llbracket \text{SEL} \rrbracket$) extracts (resp. introduces) messages from (resp. in) the global queue. The error cases for $\llbracket \text{REQ} \rrbracket$ and $\llbracket \text{SEL} \rrbracket$, namely $\llbracket \text{REQER} \rrbracket$ and $\llbracket \text{SELER} \rrbracket$, ‘skip’ the current action (removing it from the process), do not modify the queue, the router nor the state of the monitor. The error cases for $\llbracket \text{ACC} \rrbracket$ and $\llbracket \text{BRA} \rrbracket$, namely $\llbracket \text{ACCER} \rrbracket$ and $\llbracket \text{BRAER} \rrbracket$ (the latter omitted for space constraint), do not affect the process, which remains ready to perform the action, and remove the violating message from the queue.

Example 4 (ATM: a monitored network). We illustrate the monitored networks for the ATM scenario, where the routing table is defined as

$$r = a \mapsto \alpha, b \mapsto \beta, c \mapsto \gamma, s[\text{S}] \mapsto \alpha, s[\text{C}] \mapsto \beta, s[\text{A}] \mapsto \gamma$$

We consider the fragment of session where the authentication has occurred, the process of C (resp. S) is P'_C (resp. P'_S) from Example 3, and the process of A is $\mathbf{0}$.

$$\begin{aligned} \text{N}_\text{S} &= [P'_\text{S}]_\alpha \mid \text{M}_\text{S} = [s[\text{S}, \text{C}]! \text{Account}(100); P'_\text{S}]_\alpha \mid \text{M}_\text{S} \quad (\text{assuming } \text{getBalance}() = 100) \\ \text{N}_\text{C} &= [P'_\text{C}]_\beta \mid \text{M}_\text{C} = [s[\text{S}, \text{C}]? \text{Account}(x_b).P'_\text{C}]_\beta \mid \text{M}_\text{C} \\ \text{N}_\text{A} &= [\mathbf{0}]_\gamma \mid \gamma : \langle c : T_\text{A}[\text{A}] ; s[\text{A}] : \text{end} \rangle \end{aligned}$$

where $\text{M}_\text{S} = \alpha : \langle a : T_\text{S}[\text{S}] ; s[\text{S}] : \text{C}! \text{Account}(x_b : \text{int})\{x_b \geq 0\}.T'_\text{S} \rangle$ and M_C is dual.

$$\begin{aligned} \text{N}_1 &= [s[\text{S}, \text{C}]! \text{Account}(100); P'_\text{S}]_\alpha \mid \text{M}_\text{S} \mid [s[\text{S}, \text{C}]? \text{Account}(x_b).P'_\text{C}]_\beta \mid \text{M}_\text{C} \mid \text{N}_\text{A} \mid \langle r ; \emptyset \rangle \\ &\longrightarrow [P'_\text{S}]_\alpha \mid \text{M}'_\text{S} \mid [P'_\text{C}[100/x_b]]_\beta \mid \text{M}'_\text{C} \mid \text{N}_\text{A} \mid \langle r ; \emptyset \rangle \end{aligned}$$

where $\text{M}'_\text{S} = \alpha : \langle a : T_\text{S}[\text{S}] ; s[\text{S}] : T'_\text{S} \rangle$ and $\text{M}'_\text{C} = \beta : \langle b : T_\text{C}[\text{C}] ; s[\text{C}] : T'_\text{C} \rangle$

Above, $x_b \geq 0$ is satisfied since $x_b = 100$. If the server tried to communicate e.g., value -100 for x_b , the monitoring (by rule $\llbracket \text{SELER} \rrbracket$) would drop the message.

3.3 Network Satisfaction and Equivalences

Based on the formal representations of monitored networks, we now introduce the key formal tools for analysing their behaviour. First, we introduce *bisimulation* and *barbed congruence* over networks, and develop the notion of *interface*. Then we define *the satisfaction relation* $\models N : M$, used in § 3.4 to prove the properties of our framework.

Bisimulations We use M, M', \dots for a *partial network*, that is a network which does not contain a global transport, hence enabling the global observation of interactions. The labelled transition relation for processes and partial networks M is defined below.

$$\begin{aligned} (\text{REQ}) \quad & \bar{a}(s[\mathbf{r}] : T); P]_\alpha \xrightarrow{\bar{a}(s[\mathbf{r}]:T)} [\mathbf{0}]_\alpha \quad (\text{ACC}) \quad [a(y[\mathbf{r}] : T).P]_\alpha \xrightarrow{a(s[\mathbf{r}]:T)} [P[s/y]]_\alpha \\ (\text{BRA}) \quad & [s[\mathbf{r}_1, \mathbf{r}_2]? \{l_i(x_i : S_i).P_i\}_i]_\alpha \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]?l_j(v)} [P_j[v/x_j]]_\alpha \\ (\text{SEL}) \quad & [s[\mathbf{r}_1, \mathbf{r}_2]!l_j(v)]_\alpha \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l_j(v)} [\mathbf{0}]_\alpha \quad (\text{CTX}) \quad \frac{[P]_\alpha \xrightarrow{\ell} [P']_\alpha \quad \mathbf{n}(\ell) \cap \mathbf{bn}(\mathcal{E}) = \emptyset}{[\mathcal{E}(P)]_\alpha \xrightarrow{\ell} [\mathcal{E}(P')]_\alpha} \\ (\text{TAU}) \quad & \frac{M \xrightarrow{\tau} M'}{M \xrightarrow{\tau} M'} \quad (\text{RES}) \quad \frac{M \xrightarrow{\ell} M' \quad a \notin \mathbf{sbj}(\ell)}{(\nu a)M \xrightarrow{\ell} (\nu a)M'} \quad (\text{STR}) \quad \frac{M \equiv M_0 \xrightarrow{\ell} M'_0 \equiv M'}{M \xrightarrow{\ell} M'} \end{aligned}$$

In (CTX), $\mathfrak{n}(\ell)$ indicates the names occurring in ℓ while $\mathfrak{bn}(\mathcal{E})$ indicates binding \mathcal{E} induces. In (RES), $\mathfrak{sbj}(\ell)$ denotes the subject of ℓ . In (TAU) the axiom is obtained either from the reduction rules for dynamic networks given in § 2.2 (only those not involving the global transport), or from the corresponding rules for monitored networks (which have been omitted in § 3.2).

Hereafter we write \Longrightarrow for $\xrightarrow{\tau^*}$, $\xRightarrow{\ell}$ for $\Longrightarrow \xrightarrow{\ell} \Longrightarrow$, and $\xRightarrow{\hat{\ell}}$ for \Longrightarrow if $\ell = \tau$ and $\xRightarrow{\ell}$ otherwise.

Definition 1 (Bisimulation over partial networks). A binary relation \mathcal{R} over partial networks is a *weak bisimulation* when $M_1 \mathcal{R} M_2$ implies: whenever $M_1 \xrightarrow{\ell} M'_1$ such that $\mathfrak{bn}(\ell) \cap \mathfrak{fn}(M_2) = \emptyset$, we have $M_2 \xRightarrow{\hat{\ell}} M'_2$ such that $M'_1 \mathcal{R} M'_2$, and the symmetric case. We write $M_1 \approx M_2$ if (M_1, M_2) are in a weak bisimulation.

Interface We want to build a model where two different implementations of the same service are related. Bisimilarity is too strong for this aim (as shown in Example 5). We use instead a contextual congruence (barbed reduction-closed congruence [13]) \cong for networks. Intuitively, two networks are barbed-congruent when they are indistinguishable for any principal that connects to them. In this case we say they propose the same *interface* to the exterior. Formally, two networks are related with \cong when, composed with the same third network, they offer the same *barbs* (the messages to external principals in the respective global queues are on the same channels) and this property is preserved under reduction.

We say that a message m is *routed for α in N* if $N = (\nu \tilde{n})(M_0 \mid \langle r ; h \rangle)$, $m \in h$, either $m = \bar{a}\langle s[\mathbf{r}] : T \rangle$ and $r(a) = \alpha$ or $m = s[\mathbf{r}_1, \mathbf{r}_2]!l\langle e \rangle$ and $r(s[\mathbf{r}_2]) = \alpha$.

Definition 2 (Barb). We write $N \Downarrow_a$ when the global queue of N contains a message m to free a and m is routed for a principal not in N . We write $N \Downarrow_a$ if $N \xrightarrow{*} N' \Downarrow_a$.

We denote $\mathcal{P}(N)$ for a set of principals in N , $\mathcal{P}(\prod [P_i]_{\alpha_i}) = \{\alpha_1, \dots, \alpha_n\}$. We say N_1 and N_2 are *composable* when $\mathcal{P}(N_1) \cap \mathcal{P}(N_2) = \emptyset$, the union of their routing tables remains a function, and their free session names are disjoint. If N_1 and N_2 are composable, we define $N_1 \diamond N_2 = (\nu \tilde{n}_1, \tilde{n}_2)(M_1 \mid M_2 \mid \langle r_1 \cup r_2 ; h_1 \cdot h_2 \rangle)$ where $N_i = (\nu \tilde{n}_i)(M_i \mid \langle r_i ; h_i \rangle)$ ($i = 1, 2$). Notice that both equivalences are compositional, as proved in Proposition § 4.

Definition 3 (Barbed reduction-closed congruence). A relation \mathcal{R} on networks with the same principals is a *barbed r.c. congruence* [13] if the following holds: whenever $N_1 \mathcal{R} N_2$ we have: (1) for each composable N , $N \diamond N_1 \mathcal{R} N \diamond N_2$; (2) $N_1 \xrightarrow{*} N'_1$ implies $N_2 \xrightarrow{*} N'_2$ s.t. $N'_1 \mathcal{R} N'_2$ again, and the symmetric case; (3) $N_1 \Downarrow_a$ iff $N_2 \Downarrow_a$. We write $N_1 \cong N_2$ when they are related by a barbed r.c. congruence.

The following result states that composing two bisimilar partial networks with the same network – implying the same router and global transport – yields two undistinguishable networks.

Satisfaction We present a satisfaction relation for partial networks, which include local principals. If M is a partial network, $\models M : \Sigma$ s.t. $\text{dom}(\Sigma) = \mathcal{P}(M)$, means that the specification allows all outputs from the network; that the network is ready to receive all the inputs indicated by the specification; and that this is preserved by transition.

Definition 5 (Satisfaction). Let $\text{sbj}(\ell)$ denote the subject of $\ell \neq \tau$. A relation \mathcal{R} from partial networks to specifications is a *satisfaction* when $M\mathcal{R}\Sigma$ implies:

1. If $\Sigma \xrightarrow{\ell} \Sigma'$ for an input ℓ and M has an input at $\text{sbj}(\ell)$, then $M \xrightarrow{\ell} M'$ s.t. $M'\mathcal{R}\Sigma'$.
2. If $M \xrightarrow{\ell} M'$ for an output at ℓ , then $\Sigma \xrightarrow{\ell} \Sigma'$ s.t. $M'\mathcal{R}\Sigma'$.
3. If $M \xrightarrow{\tau} M'$, then $\Sigma \xrightarrow{\tau} \Sigma'$ s.t. $M'\mathcal{R}\Sigma'$ (i.e. $M'\mathcal{R}\Sigma$ since $\Sigma \xrightarrow{\tau} \Sigma$ always).

When $M\mathcal{R}\Sigma$ for a satisfaction relation \mathcal{R} , we say M *satisfies* Σ , denoted $\models M : \Sigma$. By Definition 5 and Proposition 4 we obtain:

Proposition 6. If $M_1 \cong M_2$ and $\models M_1 : \Sigma$ then $\models M_2 : \Sigma$.

3.4 Safety Assurance and Session Fidelity

In this section, we present the properties underpinning safety assurance in the proposed framework from different perspectives.

Theorem 7 shows local safety/transparency, and global safety/transparency for fully monitored networks. A network N is *fully monitored wrt* Σ when all its principals are monitored and the collection of the monitors is congruent to Σ .

Theorem 7 (Safety and Transparency).

1. **(Local Safety)** $\models [P]_\alpha \mid M : \alpha : \langle \Gamma; \Delta \rangle$ with $M = \alpha : \langle \Gamma; \Delta \rangle$.
2. **(Local Transparency)** If $\models [P]_\alpha : \alpha : \langle \Gamma; \Delta \rangle$, then $[P]_\alpha \approx ([P]_\alpha \mid M)$ with $M = \alpha : \langle \Gamma; \Delta \rangle$.
3. **(Global Safety)** If N is fully monitored w.r.t. Σ , then $\models N : \Sigma$.
4. **(Global Transparency)** Assume N and N have the same global transport $\langle r ; h \rangle$. If N is fully monitored w.r.t. Σ and $N = M \mid \langle r ; h \rangle$ is unmonitored but $\models M : \Sigma$, then we have $N \sim N$.

Local safety (7.1) states that a monitored process always behaves well with respect to the specification. Local transparency (7.2) states that a monitored process behaves as an unmonitored process when the latter is well-behaved (e.g., it is statically checked). Global safety (7.3) states that a fully monitored network behaves well with respect to the given global specification. This property is closely related to session fidelity, introduced later in Theorem 11. Global transparency (7.4) states that a monitored network and an unmonitored network have equivalent behaviour when the latter is well-behaved with respect to the same (collection of) specifications.

By Proposition 4 and (7.2), we derive Corollary 8 stating that weakly bisimilar static networks combined with the same global transport are congruent.

Corollary 8 (Local transparency). If $\models [P]_\alpha : \alpha : \langle \Gamma; \Delta \rangle$, then for any $\langle r ; h \rangle$, we have $([P]_\alpha \mid \langle r ; h \rangle) \cong ([P]_\alpha \mid M \mid \langle r ; h \rangle)$ with $M = \alpha : \langle \Gamma; \Delta \rangle$.

By Theorem 7, we can *mix* unmonitored principals with monitored principals still obtaining the desired safety assurances.

In the following, we refer to a pair $\Sigma; \langle r ; h \rangle$ of a specification and a global transport as a *configuration*. The labelled transition relation for configurations, denoted by $\xrightarrow{\ell}_g$, defined in Figure 2 in Appendix B. Here it is sufficient to notice that the transitions of a configuration define the correct behaviours (with respect to Σ) in terms of the observation of inputs and outputs from/to the global transport $\langle r ; h \rangle$. We write that a configuration $\Sigma; \langle r ; h \rangle$ is *configurationally consistent* if all of its multi-step input transition derivatives are receivable and the resulting specifications Σ is consistent.

We also use $\xrightarrow{\ell}_g$ to model globally visible transitions of networks (i.e., those locally visible transitions of a network that can be observed by its global transport). Below, we state that a message emitted by a valid output action is always receivable.

Lemma 9. Assume a network $N \equiv M \mid \langle r ; h \rangle$ conforming to $\Sigma; \langle r ; h \rangle$ which is configurationally consistent, if $N \xrightarrow{\ell}_g N'$ such that ℓ is an output and $\Sigma; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r ; h \cdot m \rangle$ then $h \cdot m$ is receivable to Σ' .

Also, we state that, as $N \equiv M \mid H$ and $\models M : \Sigma$, the satisfaction relation of M and Σ is preserved by transitions.

Lemma 10. Assume $N \equiv M \mid H$ and $\models M : \Sigma$. If $N \xrightarrow{\ell}_g N' \equiv M' \mid H'$ and $\Sigma \xrightarrow{\ell} \Sigma'$, then $\models M' : \Sigma'$.

Theorem 11 (Session Fidelity). Assume configuration $\Sigma; \langle r ; h \rangle$ is configurationally consistent, and network $N \equiv M \mid \langle r ; h \rangle$ conforms to configuration $\Sigma; \langle r ; h \rangle$. For any ℓ , whenever we have $N \xrightarrow{\ell}_g N'$ s.t. $\Sigma; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r' ; h' \rangle$, it holds that $\Sigma'; \langle r' ; h' \rangle$ is configurationally consistent and N' conforms to $\Sigma'; \langle r' ; h' \rangle$.

By session fidelity, if all session message exchanges in a monitored/unmonitored network behave well with respect to the specifications (as communications occur), then this network exactly follows the original global specifications.

4 Conclusion and future work

We proposed a new formal safety assurance framework to specify and enforce the global safety for distributed systems, through the use of static and dynamic verification. We formally proved the correctness (with respect to distributed principals) of our architectural framework through a π -calculus based theory, identified in two key properties of dynamic network: global transparency and safety. We introduced a behavioural theory over monitored networks which allows compositional reasoning over trusted and untrusted (but monitored) components.

Implementation As a part of our collaboration with the Ocean Observatories Initiative [16], our theoretical framework is currently realised by an implementation, in which each monitor supports all well-formed protocols and is automatically self-configured, via session initiation messages, for all sessions that the endpoint participates in. Our implementation of the framework automates distributed monitoring by generating FSM from the local protocol projections. In this implementation, the global protocol serves as the key abstraction that helps unify the aspects of specification, implementation and verification (both static and dynamic) of distributed application development. Our experience has shown that the specification framework can accommodate diverse practical use cases, including real-world communication patterns used in the distributed services of the OOI cyberinfrastructure [16].

Future work Our objectives include the incorporation in the implementation of more elaborate handling of error cases into monitor functionality, such as halting all local sessions or coercing to valid actions [17, 15]. In order to reach this goal, we need to combine a simplification of [4] and nested sessions [9] to handle an exception inside MPSTs. We aim to construct a simple and reliable way to raise and catch exceptions in asynchronous networks. Our work is motivated by ongoing collaborations with the Savara² and Scribble³ projects and OOI [16]. We are continuing the development of Scribble, its toolsuite and associated environments towards a full integration of sessions into the OOI infrastructure.

4.1 Related Work

Our work features a located, distributed process calculus to model monitored networks. Due to space limitations, we focus on the key differences with related work on dynamic monitoring.

The work in [12] proposes an ambient-based run-time monitoring formalism, called guardians, targeted at access control rights for network processes, and Klaim [8] advocates a hybrid (dynamic and static) approach for access control against capabilities (policies) to support static checking integrated within a dynamic access-control procedure. These works address specific forms of access control for mobility, while our more general approach aims at ensuring correct behaviour in sessions through a combination of static or run-time verification.

The work in [3] presents a monitor-based information-flow analysis in multi-party sessions. The monitors in [3] are inline (according to [5]) and control the information-flow by tagging each message with security levels. Our monitors are outline and aim at the application to distributed systems.

An informal approach to monitoring based on MPSTs, and an outline of monitors are presented in [7]. However, [7] only gives an overview of the desired properties, and requires all local processes to be dynamically verified through the protections of system monitors. In this paper, instead, we integrate statically

² <http://www.jboss.org/savara>

³ <http://www.scribble.org>

and dynamically verified local processes into one network, and formally state the properties of this combination.

In summary, compared to these related works, our contribution focuses on the enforcement of global safety, with protocols specified as multiparty session types with assertions. It also provides formalisms and theorems for decentralised runtime monitoring, targeting interaction between components written in multiple (e.g., statically and dynamically typed) programming languages.

References

1. L. Bettini et al. Global progress in dynamically interleaved multiparty sessions. In *CONCUR*, volume 5201 of *LNCS*, pages 418–433, 2008.
2. 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.
3. S. Capecchi, I. Castellani, and M. Dezani-Ciancaglini. Information flow safety in multiparty sessions. In *EXPRESS*, volume 64 of *EPTCS*, pages 16–30, 2011.
4. S. Capecchi, E. Giachino, and N. Yoshida. Global escape in multiparty session. In *FSTTCS'10*, volume 8 of *LIPICS*, pages 338–351, 2010.
5. F. Chen and G. Rosu. MOP: An Efficient and Generic Runtime Verification Framework. In *OOPSLA*, pages 569–588, 2007.
6. T.-C. Chen. *Theories for Session-based Governance for Large-Scale Distributed Systems*. PhD thesis, Queen Mary, University of London, 2013. To appear.
7. T.-C. Chen, L. Bocchi, P.-M. Deniélou, K. Honda, and N. Yoshida. Asynchronous distributed monitoring for multiparty session enforcement. In *TGC*, volume 7173 of *LNCS*, pages 25–45, 2011.
8. R. De Nicola, G. Ferrari, and R. Pugliese. Klaim: a kernel language for agents interaction and mobility. *IEEE Trans. Softw. Eng.*, 24:315–330, 1998.
9. R. Demangeon and K. Honda. Nested protocols in session types. In *CONCUR*, volume 7454 of *LNCS*, pages 272–286, 2012.
10. P.-M. Deniélou and N. Yoshida. Dynamic multirole session types. In *POPL*, pages 435–446, 2011.
11. P.-M. Deniélou and N. Yoshida. Multiparty session types meet communicating automata. In *ESOP*, volume 7211 of *LNCS*, pages 194–213, 2012.
12. G. Ferrari, E. Moggi, and R. Pugliese. Guardians for ambient-based monitoring. In *F-WAN*, pages 141–202. Elsevier, 2002.
13. K. Honda and N. Yoshida. On reduction-based process semantics. *TCS*, 151(2):437–486, 1995.
14. K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In *POPL'08*, pages 273–284. ACM, 2008.
15. J. Ligatti, L. Bauer, and D. Walker. Run-time enforcement of nonsafety policies. *ACM Trans. Inf. Syst. Secur.*, 12:19:1–19:41, 2009.
16. OOI. <http://www.oceanobservatories.org/>.
17. F. B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 3:30–50, 2000.
18. N. Yoshida, P.-M. Deniélou, A. Bejleri, and R. Hu. Parameterised multiparty session types. In *FoSSaCs'10*, volume 6014 of *LNCS*, pages 128–145, 2010.

A Safety

Theorem 7.1 (Local safety) $\models [P]_\alpha \mid \mathbf{M} : \alpha : \langle \Gamma; \Delta \rangle$ with $\mathbf{M} = \alpha : \langle \Gamma; \Delta \rangle$.

Proof. Since $[P]_\alpha$ is a partial network, by Definition 5, it is straightforward that the network $[P]_\alpha \mid \mathbf{M} :$ satisfies $\alpha : \langle \Gamma; \Delta \rangle$ due to $\mathbf{M} = \alpha : \langle \Gamma; \Delta \rangle$.

Theorem 7.2 (Local transparency (bisim)) If $\models [P]_\alpha : \alpha : \langle \Gamma; \Delta \rangle$, then $[P]_\alpha \approx ([P]_\alpha \mid \mathbf{M})$ with $\mathbf{M} = \alpha : \langle \Gamma; \Delta \rangle$.

Proof. By Theorem 7.1, we have $\models [P]_\alpha \mid \mathbf{M} : \alpha : \langle \Gamma; \Delta \rangle$ with $\mathbf{M} = \alpha : \langle \Gamma; \Delta \rangle$. Define a relation R :

$$R = \{([P]_\alpha, [P]_\alpha \mid \mathbf{M}) \mid \models [P]_\alpha : \alpha : \langle \Gamma; \Delta \rangle\}$$

Assume $([P]_\alpha, [P]_\alpha \mid \mathbf{M}) \in R$,

- for an output ℓ or τ , $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$ implies $\mathbf{M} \xrightarrow{\ell} \mathbf{M}'$ due to $\models [P]_\alpha : \mathbf{M}$, so that $[P]_\alpha \mid \mathbf{M} \xrightarrow{\ell} [P']_\alpha \mid \mathbf{M}'$;
- for an input ℓ , $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$ only when $\mathbf{M} \xrightarrow{\ell} \mathbf{M}'$, which together imply that $[P]_\alpha \mid \mathbf{M} \xrightarrow{\ell} [P']_\alpha \mid \mathbf{M}'$.

By Definition 5, we have $\models [P']_\alpha : \mathbf{M}'$, so that $([P']_\alpha, [P']_\alpha \mid \mathbf{M}') \in R$. Symmetrically,

- for an output ℓ or τ , $[P]_\alpha \mid \mathbf{M} \xrightarrow{\ell} [P']_\alpha \mid \mathbf{M}'$ implies $\mathbf{M} \xrightarrow{\ell} \mathbf{M}'$ whenever $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$;
- for an input ℓ , $[P]_\alpha \mid \mathbf{M} \xrightarrow{\ell} [P']_\alpha \mid \mathbf{M}'$ says $\mathbf{M} \xrightarrow{\ell} \mathbf{M}'$, which implies $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$.

By Definition 5, we have $\models [P']_\alpha \mid \mathbf{M}' : \mathbf{M}'$, so that $([P']_\alpha \mid \mathbf{M}', [P']_\alpha) \in R$. By Definition 1, $[P]_\alpha \approx ([P]_\alpha \mid \mathbf{M})$ with $\mathbf{M} = \alpha : \langle \Gamma; \Delta \rangle$.

Theorem 7.3 (Global safety) If \mathbf{N} is fully monitored w.r.t. Σ , then $\models \mathbf{N} : \Sigma$.

Proof. Assume \mathbf{N} is composed by monitored endpoints $[P_i]_{\alpha_i} \mid \mathbf{M}_i, i \in \{1, \dots, n\}$ and a global transport $\langle r ; h \rangle$

$$\mathbf{N} \equiv [P_1]_{\alpha_1} \mid \mathbf{M}_1 \mid \dots \mid [P_n]_{\alpha_n} \mid \mathbf{M}_n \mid \langle r ; h \rangle$$

where $\mathbf{M}_i = \alpha_i : \langle \Gamma_i; \Delta_i \rangle$ for $i = \{1, \dots, n\}$, $\Sigma = \mathbf{M}_1, \dots, \mathbf{M}_n$. Based on Theorem 7.1, for each $i \in \{1, \dots, n\}$,

$$\models [P_i]_{\alpha_i} \mid \mathbf{M}_i : \alpha_i : \langle \Gamma_i; \Delta_i \rangle$$

with $\mathbf{M}_i = \alpha_i : \langle \Gamma_i; \Delta_i \rangle$. By Definition 5 and induction, we have

$$[P_1]_{\alpha_1} \mid \mathbf{M}_1 \mid \dots \mid [P_n]_{\alpha_n} \mid \mathbf{M}_n : \alpha_1 : \langle \Gamma_1; \Delta_1 \rangle, \dots, \alpha_n : \langle \Gamma_n; \Delta_n \rangle$$

so that $\models \mathbf{N} : \Sigma$.

As the rules of reduction relation for dynamic networks, introduced in Section 2.2, define the *invisible* actions for proving the properties of *global transparency* (Theorem 7.4) and *session fidelity* (Theorem 11), we introduce the LTS of dynamic networks in Figure 1, where the actions are *observable* through observing the dynamics of the global transport h . The notation of global observable transition $\xrightarrow{\ell}_g$ is used to denote globally observable action ℓ . For unmonitored

$$\begin{array}{l}
\{\text{REQ}\} \quad \langle r ; h \rangle \xrightarrow{\bar{a}\langle s[\mathbf{x}]:T \rangle}_g \langle r ; h \cdot \bar{a}\langle s[\mathbf{x}]:T \rangle \rangle \\
\{\text{ACC}\} \quad \langle r ; \bar{a}\langle s[\mathbf{x}]:T \rangle \cdot h \rangle \xrightarrow{a\langle s[\mathbf{x}]:T \rangle}_g \langle r ; h \rangle \\
\{\text{SEL}\} \quad \langle r ; h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle}_g \langle r ; h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle \rangle \\
\{\text{BRA}\} \quad \langle r ; s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle \cdot h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]?l\langle v \rangle}_g \langle r ; h \rangle \\
\{\text{IN}\} \quad \frac{\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle \quad \ell \text{ input} \quad \text{dest}(\ell, r) \in \mathcal{P}(M)}{M \mid \langle r ; h \rangle \xrightarrow{\ell}_g M \mid \langle r' ; h' \rangle} \\
\{\text{OUT}\} \quad \frac{\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle \quad \ell \text{ output} \quad \text{dest}(\ell, r) \notin \mathcal{P}(M)}{M \mid \langle r ; h \rangle \xrightarrow{\ell}_g M \mid \langle r' ; h' \rangle} \\
\{\text{TAU}\} \quad \frac{N \longrightarrow N'}{N \xrightarrow{\tau} N'} \quad \{\text{RES}\} \quad \frac{N \xrightarrow{\ell} N' \quad a \notin \text{sbj}(\ell)}{(va)N \xrightarrow{\ell \setminus a} (va)N'} \\
\{\text{STR}\} \quad \frac{N \equiv N_0 \xrightarrow{\ell} N'_0 \equiv M'}{N \xrightarrow{\ell} N'} \quad \{\text{PAR}\} \quad \frac{N_1 \xrightarrow{\ell}_g N'_1 \quad \text{bn}(\ell) \cap \text{fn}(N_2) = \emptyset \quad \text{dest}(\ell) \notin \mathcal{P}(\cdot)N_2}{N_1 \parallel N_2 \xrightarrow{\ell}_g N'_1 \mid N_2}
\end{array}$$

Fig. 1. LTS for dynamic networks

networks, as $N \xrightarrow{\ell}_g N'$, it means $\exists [P]_\alpha \in N$, $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$ (i.e. locally visible) such that $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$ (i.e. globally visible). Similarly, for monitored networks, $N \xrightarrow{\ell}_g N'$ means $\exists [P]_\alpha \mid M \in N$, $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$ and $M \xrightarrow{\ell} M'$ (i.e. locally visible) such that $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$ (i.e. globally visible). As a message leaves the global queue, there should exist a local process receiving it as an input. Rules $\{\text{ACC}\}$ and $\{\text{BRA}\}$ are generalised as rule $\{\text{IN}\}$, and rules $\{\text{REQ}\}$ and $\{\text{SEL}\}$ are generalised as rule $\{\text{OUT}\}$. As we explained above, the visible actions in these rules are observable in the global queue. Rule $\{\text{OUT}\}$ means there is a local process in M outputting a message to the global queue, thus the destination of action ℓ should not be inside M but outside it. Similarly, rule $\{\text{IN}\}$ means there is a local process in M going to receive a message from the global queue, thus the destination of action ℓ should be inside M . Rule $\{\text{TAU}\}$ summarise the reduction rules defined in Section 2.2. Rule $\{\text{RES}\}$ and $\{\text{STR}\}$ are standard. Rule $\{\text{PAR}\}$ says that, the bound names of action ℓ should not be any free name appearig in network N_2 , and it should not be absorbed by any process in network N_2 (i.e. its destination is not in N_2).

Theorem 7.4 (Global transparency) Assum \mathbf{N} and N have the same global transport $\langle r ; h \rangle$. If \mathbf{N} is fully monitored w.r.t. Σ and $N = M \mid \langle r ; h \rangle$ is unmonitored but $\models M : \Sigma$, then we have $\mathbf{N} \sim N$.

Proof. Define a relation R :

$$R = \{ \mathbf{N}, N \mid N = M \mid \langle r ; h \rangle \text{ and } \models M : \Sigma \}$$

We prove that R is a standard strong bisimilar relation over $\xrightarrow{\ell}_{\mathbf{g}}$. Note that, $M : \Sigma$ means $\forall [P]_{\alpha_i} \in M$, we have $\alpha_i : \langle \Gamma_i ; \Delta_i \rangle \in \Sigma$ and $\models [P]_{\alpha_i} : \alpha_i : \langle \Gamma_i ; \Delta_i \rangle$.

1. As $\mathbf{N} \xrightarrow{\ell}_{\mathbf{g}} \mathbf{N}'$, it implies $\exists [P]_{\alpha_j} \mid M_j \in \mathbf{N}$, $[P]_{\alpha_j} \xrightarrow{\ell} [P']_{\alpha_j}$ and $M_j \xrightarrow{\ell} M'_j$ such that $\langle r ; h \rangle \xrightarrow{\ell}_{\mathbf{g}} \langle r' ; h' \rangle$, and other monitored processes in \mathbf{N} are not affected. When ℓ is an input, by Definition 5, since $\models M : \Sigma$, we should have $[P]_{\alpha_j} \xrightarrow{\ell} [P']_{\alpha_j}$; when ℓ is an output or a τ action, by Definition 5, the transition of $[P]_{\alpha_j} \xrightarrow{\ell} [P']_{\alpha_j}$ is able to take place. Both cases lead to $M \xrightarrow{\ell} M'$ and $\langle r ; h \rangle \xrightarrow{\ell}_{\mathbf{g}} \langle r' ; h' \rangle$ so that $N = M \mid \langle r ; h \rangle \xrightarrow{\ell}_{\mathbf{g}} M' \mid \langle r' ; h' \rangle = N'$, and $\models [P']_{\alpha_j} : \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle$ by Definition 5. $\alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle$ is the resulting new configuration of α_j in Σ . Other specifications $\{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I \setminus \{j\}} \in \Sigma$ are not affected. Let $\Sigma' = \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle, \{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I \setminus \{j\}}$. Therefore, for the resulting new network $N' = M' \mid \langle r' ; h' \rangle$, we have $\models M' : \Sigma'$. Thus we have $(\mathbf{N}', N') \in R$.
2. For the symmetric case, as $N \xrightarrow{\ell}_{\mathbf{g}} N'$, it implies $\exists [P]_{\alpha_j} \in N$, $[P]_{\alpha_j} \xrightarrow{\ell} [P']_{\alpha_j}$ such that $\langle r ; h \rangle \xrightarrow{\ell}_{\mathbf{g}} \langle r' ; h' \rangle$ and other processes in N are not affected. Since $\models M : \Sigma$, without loss of generality, let $M_j = \alpha_j : \langle \Gamma_j ; \Delta_j \rangle$, then we have, for any ℓ , $[P]_{\alpha_j} \mid M_j \xrightarrow{\ell} [P']_{\alpha_j} \mid M'_j$, where $M'_j = \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle$. It makes $\langle r ; h \rangle \xrightarrow{\ell}_{\mathbf{g}} \langle r' ; h' \rangle$, so that $N \xrightarrow{\ell}_{\mathbf{g}} N'$. Since N' is a fully monitored network, its static part (i.e. the part when the global transport is taken off from N'), say $[P]_{\alpha_i} \mid \{M_i\}_{i \in I}$ where $\{M_i\}_{i \in I} = \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle, \{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I \setminus \{j\}}$, $\models [P]_{\alpha_i} \mid \{M_i\}_{i \in I} : \Sigma'$ where $\Sigma' = \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle, \{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I \setminus \{j\}}$. Thus we have $(N', N') \in R$.

B Session Fidelity

To prove the property of *session fidelity*, we define the LTS of configurations in Figure 2. All rules are straightforward from the LTS of specifications and the one of networks.

Definition 12 (Configuration). A configuration is denoted by $\Phi = \Sigma ; \langle r ; h \rangle$, in which the group of monitors correspond to h . In other words, all messages corresponding to the actions guarded by Σ are in h .

A Φ thus guides and captures the behaviours in the network. Let $\mathcal{P}(\Phi)$ be the set of principals involving in Φ . We define the composition of configurations as follows.

$$\begin{array}{c}
\text{[Req]} \quad \frac{\Sigma \xrightarrow{\bar{a}\langle s[\mathbf{r}] : T \rangle} \Sigma'}{\Sigma ; \langle r ; h \rangle \xrightarrow{\bar{a}\langle s[\mathbf{r}] : T \rangle}_{\mathbf{g}} \Sigma' ; \langle r ; h \cdot \bar{a}\langle s[\mathbf{r}] : T \rangle \rangle} \\
\text{[Acc]} \quad \frac{\alpha : \langle \Gamma, a : \mathbf{I}(T[\mathbf{r}]); \Delta \rangle \in \Sigma \quad \Sigma \xrightarrow{a\langle s[\mathbf{r}] : T \rangle} \Sigma'}{\Sigma ; \langle r ; \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot h \rangle \xrightarrow{a\langle s[\mathbf{r}] : T \rangle}_{\mathbf{g}} \Sigma' ; \langle r, s[\mathbf{r}] \mapsto \alpha ; h \rangle} \\
\text{[Sel]} \quad \frac{\Sigma \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle} \Sigma'}{\Sigma ; \langle r ; h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle}_{\mathbf{g}} \Sigma' ; \langle r ; h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle \rangle} \\
\text{[Bra]} \quad \frac{\Sigma \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]?l\langle v \rangle} \Sigma'}{\Sigma ; \langle r ; s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle \cdot h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]?l\langle v \rangle}_{\mathbf{g}} \Sigma' ; \langle r ; h \rangle} \\
\text{[Par]} \quad \frac{\Phi_1 \xrightarrow{\ell}_{\mathbf{g}} \Phi_2}{\Phi_1 \odot \Phi_3 \xrightarrow{\ell}_{\mathbf{g}} \Phi_2 \odot \Phi_3} \quad \text{[Tau]} \quad \frac{\Sigma \xrightarrow{\tau} \Sigma'}{\Sigma ; \langle r ; h \rangle \xrightarrow{\tau}_{\mathbf{g}} \Sigma ; \langle r ; h \rangle}
\end{array}$$

Fig. 2. Labelled transition relation for configurations

Definition 13 (Parallel composition of configurations). Assume $\Phi_1 = \Sigma_1 ; \langle r_1 ; h_1 \rangle$ and $\Phi_2 = \Sigma_2 ; \langle r_2 ; h_2 \rangle$, we say Φ_1 and Φ_2 are *composable* whenever $\mathcal{P}(\Phi_1) \cap \mathcal{P}(\Phi_2) = \emptyset$ and the union of their routing tables remains a function. If Φ_1 and Φ_2 are composable, we define the composition of Φ_1 and Φ_2 as: $\Phi_1 \odot \Phi_2 = \Sigma_1, \Sigma_2 ; \langle r_1 \cup r_2 ; h_1 \cdot h_2 \rangle$.

The behaviours of each principal in a network are guided by the Σ (specification), and are observed by the $\langle r ; h \rangle$ (global transport). Except rules [Acc] and [Par], all rules are straightforward from the LTS of specifications (defined in Section 3.1) and the one of dynamic networks (Figure 1).

1. Rule [Acc] indicates that, only when the invitation has been (internally) accepted by a principal in the network, the routing information registers $s[\mathbf{r}] \mapsto \alpha$. When we observe the global transport (externally), we only observe that an invitation is moved out from the global queue (which implies that it has been accepted). However, we do not know *who* accepts it. Only Σ tells which principal accepts this invitation, so that we can register it in the routing information using α .
2. Rule Par says if Φ_1 and Φ_3 are composable (Definition 13), after Φ_1 becomes as Φ_2 , they are still composable.

Definition 14 (Consistency, Coherence). $\Sigma = \{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I}$ is *consistent* when

1. there is one and only one i such that $\Gamma_i \vdash a : \mathbf{I}(T[\mathbf{r}])$, and
2. as long as $a : \mathbf{0}(T[\mathbf{r}])$ exists in some Γ_i , $\exists \Gamma_j$ such that $a : \mathbf{I}(T[\mathbf{r}]) \in \Gamma_j$; and
3. for any s appearing in any Δ_j , if $\{s[\mathbf{r}_k] : T_k\}_{1 \leq k \leq n}$ is a collection appeared in $\{\Delta_i\}_{i \in I}$, there exists well-formed G such that $\text{roles}(G) = \{\mathbf{r}_1, \dots, \mathbf{r}_n\}$ and $G \upharpoonright \mathbf{r}_i = T_i$.

Two specifications Σ_1 and Σ_2 are *coherent* when their union is a consistent specification.

Definition 15 (Routing table). We define $\text{route}(\Sigma)$, the routing table derived from Σ , as follows:

$$\begin{aligned}\text{route}(\alpha : \langle \Gamma; \Delta, s[\mathbf{r}] : T \rangle, \Sigma) &= s[\mathbf{r}] \mapsto \alpha, \text{route}(\alpha : \langle \Gamma; \Delta \rangle, \Sigma) \\ \text{route}(\alpha : \langle \Gamma, a : \mathbf{I}(T[\mathbf{r}]); \Delta \rangle, \Sigma) &= a \mapsto \alpha, \text{route}(\alpha : \langle \Gamma; \Delta \rangle, \Sigma) \\ \text{route}(\alpha : \langle \Gamma, a : \mathbf{O}(T[\mathbf{r}]); \Delta \rangle, \Sigma) &= \text{route}(\alpha : \langle \Gamma; \Delta \rangle, \Sigma)\end{aligned}$$

for $\text{route}(\alpha : \langle \Gamma, a : \mathbf{O}(T[\mathbf{r}]) \rangle)$, because $a : \mathbf{O}(T[\mathbf{r}])$ implies that $a : \mathbf{I}(T[\mathbf{r}])$ should exist in the network, routing table should have contained the routing information for a .

The theorem of session fidelity states that, whenever a network conforms to specifications, i.e., its all local processes (static network) conform to specifications, all of its derivatives conform to specifications. In the follows, we firstly formally define *receivability*, *consistency* and *conformance* based on LTS of configurations and dynamic networks.

Definition 16 (Receivable Configuration). Define $\Sigma; \langle r ; h \rangle$ is *receivable* by the following induction:

1. If h is empty then $\Sigma; \langle r ; h \rangle$ is receivable.
2. If $h \equiv m \cdot h'$, then $\Sigma; \langle r ; h \rangle$ is receivable when we have $\Sigma; \langle r ; m \cdot h' \rangle \xrightarrow{\ell}_g \Sigma'; \langle r' ; h' \rangle$, where ℓ corresponding to m , and $\Sigma'; \langle r' ; h' \rangle$ is receivable.

A configuration $\Sigma; \langle r ; h \rangle$ is *configurationally consistent* if all of its multi-step global input transition derivatives are receivable and the resulting specifications Σ is consistent. The consistency of specifications is defined in Definition 14.

Definition 17 (Configurational Consistency). A configuration $\Phi = \Sigma; \langle r ; h \rangle$ is *configurationally consistent* whenever

1. If h is empty and Σ is consistent, or
2. h is not empty, the sequence of messages in h are receivable to Σ , and after receiving all messages in h with $\Sigma \xrightarrow{\ell_1 \dots \ell_n} \Sigma'$, where $\ell_i, i = \{1, \dots, n\}$ are inputs and, $\forall m \in h, \exists \ell \in \ell_1 \dots \ell_n$ such that ℓ corresponds to m , we have Σ' is consistent.

In other words, $\Sigma; \langle r ; h \rangle$ is configurationally consistent if, in each of its derivatives, all messages in the transport can be “received” by some monitors in Σ and, after absorbing all these messages, the resulting Σ' is still consistent.

Definition 18 (Conformance to a Configuration). Assume a network $N \equiv M \mid \langle r ; h \rangle$ is given. Define N conforms to $\Sigma; \langle r ; h \rangle$ when:

1. h is empty, $\models M : \Sigma$ and Σ is consistent, or
2. h is not empty, and the following conditions hold
 - (a) $\models M : \Sigma$,
 - (b) all messages in h are receivable to M , and

(c) as $\Sigma; \langle r ; h \rangle \xrightarrow{\ell_1 \dots \ell_n}_g \Sigma'; \langle r' ; \emptyset \rangle$ so that $M \mid h \xrightarrow{\ell_1 \dots \ell_n}_g M' \mid \emptyset$ where each $\ell_i, i = \{1, \dots, n\}$ is an input, Σ' is consistent.

The following session fidelity theorem states: assume network $N \equiv M \mid \langle r ; \emptyset \rangle$ is given, suppose that M satisfies Σ . If Σ is consistent and if $r = \text{route}(\Sigma)$, then we say N conforms to Σ . If this holds, then, with the messages which N exchanges follow the specification, the dynamics of the network witnesses the validity of specifications. *In the follows, we always assume Σ is consistent, unless otherwise stated.*

Theorem 19 (Session Fidelity). Assume configuration $\Sigma; \langle r ; h \rangle$ is configurationally consistent. and network $N \equiv M \mid \langle r ; h \rangle$ conforms to configuration $\Sigma; \langle r ; h \rangle$. Then for any ℓ , whenever we have $N \xrightarrow{\ell}_g N'$ s.t. $\Sigma; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r' ; h' \rangle$, it holds that $\Sigma'; \langle r' ; h' \rangle$ is configurationally consistent and that N' conforms to $\Sigma'; \langle r' ; h' \rangle$.

Before proving the property of session fidelity, we first prove the following lemmas.

Lemma 9 Assume a network $N \equiv M \mid \langle r ; h \rangle$ conforming to $\Sigma; \langle r ; h \rangle$ which is configurationally consistent, if $N \xrightarrow{\ell}_g N'$ such that ℓ is an output and $\Sigma; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r ; h \cdot m \rangle$ then $h \cdot m$ is receivable to Σ' .

Proof. When $\ell = \bar{a}\langle s[\mathbf{r}] : T \rangle$, since Σ is consistent, by Definitions 14, there exists $a : \mathbf{I}(T[\mathbf{r}])$ in some Γ of Σ . Because ℓ does not affect the existence of $a : \mathbf{I}(T[\mathbf{r}])$, it remains in Γ of Σ' , thus invitation $m = \bar{a}\langle s[\mathbf{r}] : T \rangle$ is receivable to Σ' .

Let $\alpha_i = \langle \Gamma_i, \Delta_i \rangle$. When $\ell = s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle$, by Definitions 14 and 18, since $\models M : \Sigma$ and Σ is consistent, $\exists \alpha_s, \alpha_r \in \Sigma$, $\exists G$ is well-formed and s obeys to G ,

$$G = \mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{l_i(x_i : (T[\mathbf{p}])_i)\{A_i\}.G_i\}_{i \in I}$$

such that

$$\begin{aligned} \Delta_s(s[\mathbf{r}_1]) &= G \upharpoonright \mathbf{r}_1 = \mathbf{r}_2! \{l_i(x_i : (T[\mathbf{p}])_i)\{A_i\}.G_i \upharpoonright \mathbf{r}_1\}_{i \in I} \\ \Delta_r(s[\mathbf{r}_2]) &= G \upharpoonright \mathbf{r}_2 = \mathbf{r}_1? \{l_i(x_i : (T[\mathbf{p}])_i)\{A'_i\}.G_i \upharpoonright \mathbf{r}_2\}_{i \in I} \end{aligned} \quad (1)$$

As action $s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle$ fires, Equation 1 changes to

$$\begin{aligned} \Delta_s(s[\mathbf{r}_1]) &= G_j \upharpoonright \mathbf{r}_1 \\ \Delta_r(s[\mathbf{r}_2]) &= G \upharpoonright \mathbf{r}_2 = \mathbf{r}_1? \{l_i(x_i : (T[\mathbf{p}])_i)\{A'_i\}.G_i \upharpoonright \mathbf{r}_2\}_{i \in I} \end{aligned} \quad (2)$$

the receiving capability of $\mathbf{r}_1?$ still remains in $\Delta_r(s[\mathbf{r}_2])$, where $\alpha_r \in \Sigma'$, thus $m = s[\mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle]$ is receivable to Σ' .

As $N \equiv M \mid H$ and $\models M : \Sigma$, the satisfaction relation of M and Σ remains whenever action takes place.

Lemma 10 Assume $N \equiv M \mid H$ and $\models M : \Sigma$. If $N \xrightarrow{\ell}_g N' \equiv M' \mid H'$ and $\Sigma \xrightarrow{\ell} \Sigma'$, then $\models M' : \Sigma'$.

Proof. Directly from Definition 5.

Now we prove session fidelity:

Proof. Assume N conforms to $\Sigma; \langle r ; h \rangle$, which is configurationally consistent. We prove the statement by inspection of each case.

(Sel) Let $\ell = s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle$, $N \xrightarrow{\ell}_g N'$ and $\Sigma; \langle \text{route}(\Sigma) ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r ; h \cdot m \rangle$, where $m = s\langle \mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle \rangle$.

Then $r = \text{route}(\Sigma) = \text{route}(\Sigma')$ because there is no change to the elements in Σ or to the routing table.

Since Σ allows ℓ , and Σ is consistent, $\exists \alpha_r, \alpha_s \in \Sigma$, $\exists G$ is well-formed,

$$G = \mathbf{r}_1 \rightarrow \mathbf{r}_2 \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I},$$

such that

$$\begin{aligned} \Delta_s(s[\mathbf{r}_1]) &= G \upharpoonright \mathbf{r}_1 = \mathbf{r}_2! \{l_i(x_i : S_i)\{A_i\}.G_i \upharpoonright \mathbf{r}_1\}_{i \in I}, \\ \Delta_r(s[\mathbf{r}_2]) &= G \upharpoonright \mathbf{r}_2 = \mathbf{r}_1? \{l_i(x_i : S_i)\{A'_i\}.G_i \upharpoonright \mathbf{r}_2\}_{i \in I}. \end{aligned}$$

$\Sigma \xrightarrow{\ell} \Sigma'$ implies Σ' has

$$\begin{aligned} s[\mathbf{r}_1] &= G_j \upharpoonright \mathbf{r}_1, \\ s[\mathbf{r}_2] &= \mathbf{r}_1? \{l_i(x_i : S_i)\{A'_i\}.G_i \upharpoonright \mathbf{r}_2\}_{i \in I}. \end{aligned}$$

Consider (case 1: h is empty). By Lemma 9, after receiving m , say $\Sigma' \xrightarrow{\ell} \Sigma''$, Σ'' has $s[\mathbf{r}_1] = G_j \upharpoonright \mathbf{r}_1$ and $s[\mathbf{r}_2] = G_j \upharpoonright \mathbf{r}_2$, Σ'' is thus consistent by Definition 14. By Definition 17, $\Sigma'; \langle r ; m \rangle$ is configurationally consistent, and $\models M' : \Sigma'$ by Lemma 10, thus N' conforms to $\Sigma'; \langle r ; h \cdot m \rangle$.

Consider (case 2: h is not empty). Since $\Sigma; \langle r ; h \rangle$ is configurationally consistent, again, by Lemma 9, after receiving messages in h (but not m), say $\Sigma' \xrightarrow{\ell_0 \dots \ell_n} \Sigma'_1$, where every action in $\ell_0 \dots \ell_n$ corresponds to each message in h , we have $\Sigma'_1; \langle r' ; m \rangle$ is configurationally consistent. After Σ'_1 receives m , say $\Sigma'_1 \xrightarrow{s[\mathbf{p}_1, \mathbf{p}_2]?l\langle v \rangle} \Sigma''$, where $s[\mathbf{p}_1, \mathbf{p}_2]?l\langle v \rangle$ is dual to ℓ , with the same reasoning above, Σ'' has $s[\mathbf{r}_1] = G'_j \upharpoonright \mathbf{r}_1$ and $s[\mathbf{r}_2] = G'_j \upharpoonright \mathbf{r}_2$, so that Σ'' is consistent. By Definition 17, $\Sigma'; \langle r ; h \cdot m \rangle$ is configurationally consistent, and $\models M' : \Sigma'$ by Lemma 10, thus N' conforms to $\Sigma'; \langle r ; h \cdot m \rangle$.

(Bra) Let $\ell = s[\mathbf{r}_1, \mathbf{r}_2]?l_j\langle v \rangle$, $N \xrightarrow{\ell}_g N'$ and N conforms to $\Sigma; \langle \text{route}(\Sigma) ; h \rangle$.

Consider (case 1: h is empty). Since $\Sigma; \langle \text{route}(\Sigma) ; \emptyset \rangle \not\rightarrow_g^\ell$, so this case never happens.

Consider (case 2: h is not empty). When h is not empty. $N \xrightarrow{g}^\ell N'$ and

$$\Sigma; \langle \text{route}(\Sigma) ; h \rangle \xrightarrow{g}^\ell \Sigma'; \langle r ; h/m \rangle,$$

where h/m means taking off message m from h , where $m = s\langle \mathbf{r}_1, \mathbf{r}_2, l_j \langle v \rangle \rangle$

We have $r = \text{route}(\Sigma) = \text{route}(\Sigma')$ because there is no change to the elements in Σ or to the routing table. By Definition 17, after receiving all messages in H , Σ is consistent, thus Σ' , which has received message m is consistent after receiving all messages in h/m . By Lemma 10, we have $\models M' : \Sigma'$ thus N' conforms to $\Sigma'; \langle r ; h/m \rangle$.

(Req) Let $\ell = \bar{a}\langle s[\mathbf{r}] : T \rangle$. $N \xrightarrow{g}^\ell N'$ and

$$\Sigma; \langle \text{route}(\Sigma) ; h \rangle \xrightarrow{g}^\ell \Sigma'; \langle r ; h \cdot m \rangle,$$

where $m = \bar{a}\langle s[\mathbf{r}] : T \rangle$. Then $r = \text{route}(\Sigma) = \text{route}(\Sigma')$ because, by Definition 15, nothing new is registered to the routing table.

Since Σ allows ℓ and Σ is consistent, by Definition 14, $\exists \Gamma_i, \Gamma_j \in \Sigma$ such that $a : \mathbf{I}(T[\mathbf{r}]) \in \Gamma_i$ and $a : \mathbf{O}(T[\mathbf{r}]) \in \Gamma_j$. After $\Sigma \xrightarrow{g}^\ell \Sigma'$, by rule [REQ] in the LTS of specifications, $a : \mathbf{I}(T[\mathbf{r}])$ remains in Γ'_i , $a : \mathbf{O}(T[\mathbf{r}])$ remains in Γ'_j , and thus they both remain in Σ' .

Consider (case 1: h is empty): By Lemma 9, after receiving m , say $\Sigma' \xrightarrow{a\langle s[\mathbf{r}]:T \rangle} \Sigma''$, both $a : \mathbf{I}(T[\mathbf{r}])$ and $a : \mathbf{O}(T[\mathbf{r}])$ remain in Σ'' , satisfying Definition 14, so that $\Sigma'; \langle r ; m \rangle$ is configurationally consistent. By Lemma 10, we have $\models M' : \Sigma'$, thus N' conforms to $\Sigma'; \langle r ; h \cdot m \rangle$.

Consider (case 2: h is not empty). The proof is similar to the one in (Sel) and omitted.

(Acc) Let $\ell = a\langle s[\mathbf{r}] : T \rangle$.

Consider (case 1: h is empty). Since $\Sigma; \langle \text{route}(\Sigma) ; \emptyset \rangle \not\rightarrow_g^\ell$, this case never happens.

Consider (case 2: h is not empty). If $N \xrightarrow{g}^\ell N'$ and

$$\Sigma; \langle \text{route}(\Sigma) ; h \rangle \xrightarrow{g}^\ell \Sigma'; \langle r' ; h/m \rangle,$$

where $m = \bar{a}\langle s[\mathbf{r}] : T \rangle$. Since there exists $\Delta \in \Sigma'$ s.t. $s[\mathbf{r}] \in \Delta$, by Definition 15, $r' = \text{route}(\Sigma)$, $s[\mathbf{r}] \mapsto \alpha = \text{route}(\Sigma')$.

For the same reasoning in (Bra), we have $\Sigma'; \langle r ; h/m \rangle$ is configurationally consistent. By Lemma 10, we have $\models M' : \Sigma'$ thus N' conforms to $\Sigma'; \langle r ; h/m \rangle$.

The proof for other cases are trivial.

C Satisfaction

Proposition 4 [Congruency] If $M_1 \approx M_2$, then

- (1) $M_1|M \approx M_2|M$ for each composable partial M ; and
- (2) $M_1|N \cong M_2|N$ for each composable N .

Proof. For (1) we show that the relation

$$\mathcal{R} = \{(M_1|M, M_2|M) \mid M_1 \approx M_2, M \text{ composable with } M_1 \text{ and } M_2\}$$

is a bisimulation. Suppose $(M_1|M)\mathcal{R}(M_2|M)$ and $M_1|M \xrightarrow{\ell} M^1$. We discuss the shape of M^1 :

- If $M^1 = M'_1|M$, it means that $M_1 \xrightarrow{\ell} M'_1$. By definition of \mathcal{R} , $M_2 \xRightarrow{\hat{\ell}} M'_2$ and $M'_1 \approx M'_2$, we conclude.
- If $M^1 = M_1|M'$, it means that $M \xrightarrow{\ell} M'$. It is easy to conclude.

By examining the reduction rule associated to parallel composition, we observe no reduction is induced through interactions between two networks. Hence we have covered all cases. The symmetric case (when $M_2|M \xrightarrow{\ell} M^2$) is easy.

To prove (2) we proceed by showing that

$$\mathcal{R} = \{((\nu\tilde{n})(M_1|N), (\nu\tilde{n})(M_2|N)) \mid M_1 \approx M_2, N \text{ comp. with } M_1, M_2\}$$

is a barbed congruence. First, this is clearly a congruence since it is closed under composition. Second, for (1), we take a composable N' . We have $N' \diamond (M_i|N) = M_i|(N' \diamond N)$. We use the definition of \mathcal{R} to conclude. For (2), assume $M_1|N \longrightarrow N^1$.

- If $N^1 = M_1|N'$, meaning that $N \longrightarrow N'$. We use the definition of \mathcal{R} to conclude.
- If $N^1 = M'_1|N'$, meaning that $N = M_0|\langle r ; \ell \cdot H \rangle$, $N = M'_0|\langle r ; H \rangle$ and $M_1 \xrightarrow{\ell} M'_1$. We deduce $N^2 = M'_2|N'$, with $N = M_0|\langle r ; \ell H \rangle$, $N = M'_0|\langle r ; H \rangle$ and $M_2 \xrightarrow{\ell} M'_2$. We use the definition of \mathcal{R} to conclude.
- If the reduction is induced by interaction between M_1 and N , then M_2 has the corresponding action, hence we can reason in the same way, hence done.

For (2), we suppose that $(M_1|N) \Downarrow_{\ell}$. Two cases can occur:

- Either $N \Downarrow_{\ell}$ and it follows directly that $(M_2|N) \Downarrow_{\ell}$.
- or $M_1 \xrightarrow{\ell} M'_1$ and by definition of \mathcal{R} , $M_2 \xRightarrow{\ell} M'_2$, meaning that $(M_2|N) \Downarrow_a$.

The symmetric case is similar.

By definition this shows $\approx \subseteq \cong$.

The two satisfactions are related by the following cut-rule-like composition principle, which enables a composition of a new partial network to a full network, as Proposition 23.

Proposition 20 (Determinism). $\Sigma \xrightarrow{\ell} \Sigma_1$ and $\Sigma \xrightarrow{\ell} \Sigma_2$ imply $\Sigma_1 = \Sigma_2$.

Proposition 20 does *not* mean a behaviour satisfying a specification is deterministic. The determinism is essential for our dynamic verification to predictably guarantee safety properties.

Proof. Suppose $\Sigma \xrightarrow{\ell} \Sigma_1$ and $\Sigma \xrightarrow{\ell} \Sigma_2$. We discuss the nature of ℓ .

- If $\ell = \bar{a}\langle s[\mathbf{r}] : T \rangle$. We deduce $\Sigma = \Sigma^0, \alpha : \langle \Gamma, a : \mathbf{I}(T[\mathbf{r}]); \Delta \rangle$ and $\Sigma_1 = \Sigma^0, \alpha : \langle a : \mathbf{I}(T[\mathbf{r}]), \Gamma; \Delta \rangle$. The definition of Σ ensures that $\alpha \notin \Sigma^0$. We deduce $\Sigma_2 = \Sigma^0, \alpha : \langle a : \mathbf{I}(T[\mathbf{r}]), \Gamma; \Delta \rangle$.
- If $\ell = s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle$. We deduce $\Sigma = \Sigma^0, \alpha : \langle \Gamma; \Delta, s[\mathbf{r}_1] : T \rangle$ and $\Sigma_1 = \Sigma_0, \alpha : \langle \Gamma; \Delta, s[\mathbf{r}_1] : T'_j\{v/x_j\} \rangle$. The definition of Σ ensures that $\alpha \notin \Sigma_0$.
- The other cases are similar or straightforward.

Hereafter we denote the set of roles in G with $\text{roles}(G)$.

The following states that if a partial network satisfies a specification, then we can add a global transport to it to obtain a full network satisfying any coherent specification.

Proposition 21 (Completion of partial network). Let M_0 be a partial network s.t. $\models M_0 : \Sigma_1$. If Σ_2 is coherent with Σ_1 , then $\models \Sigma_2 \triangleright (M_0 | \langle \text{route}(\Sigma_1); \emptyset \rangle)$.

Proof. In the proof we use a “processisation” of a queue in order to identify, from an external point of view, the processes

$$(M_1 | \langle \text{route}(M_1); \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot H \rangle)$$

and

$$(M_1 | \bar{a}\langle s[\mathbf{r}] : T \rangle . | \langle \text{route}(M_1); H \rangle).$$

This operation is defined as follows:

$$\begin{aligned} \text{Pss}((M_1 | \langle \text{route}(M_1); \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot H \rangle)) &= \\ \text{Pss}((M_1 | \bar{a}\langle s[\mathbf{r}] : T \rangle) | \langle \text{route}(M_1); H \rangle)); & \\ \text{Pss}((M_1 | \langle \text{route}(M_1); s[\mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle] \cdot H \rangle)) &= \\ \text{Pss}((M_1 | s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle; P | \langle \text{route}(M_1); H \rangle)); & \\ \text{Pss}((M_1 | \langle \text{route}(N_1); \emptyset \rangle)) &= \\ (M_1 | \langle \text{route}(M_1); \emptyset \rangle) & \end{aligned}$$

We prove that the relation

$$\mathcal{R} = \{((M_1 | \langle \text{route}(M_1); H \rangle), \Sigma_2), \text{ for all } M_1, \Sigma_2 \text{ such that } \models \text{Pss}(M_1) : \Sigma_1 \text{ for some } \Sigma_1 \text{ coherent with } \Sigma_2\}$$

is an external satisfaction relation. We have the following cases:

- If $\Sigma_2 \xrightarrow{\ell} \Sigma'_2$ then, as Σ_1 is coherent with Σ_2 , we have $\Sigma_1 \xrightarrow{\bar{\ell}} \Sigma'_1$ with Σ'_1 coherent with Σ'_2 . By the internal satisfaction relation, we have $\text{Pss}(M_1 | \langle \text{route}(M_1); H \rangle) \xrightarrow{\bar{\ell}} \text{Pss}(M'_1 | \langle \text{route}(M_1); H' \rangle)$ and $\models \text{Pss}(M_1 | \langle \text{route}(M_1); H' \rangle) : \Sigma'_1$. By definition of \mathcal{R} , we conclude.
- If $(M_1 | \langle \text{route}(M_1); H \rangle) \xrightarrow{\ell} N'$, meaning that $N' = (M_1 | \langle \text{route}(M_1); H' \rangle)$, we have

$$\text{Pss}(M_1 | \langle \text{route}(M_1); H \rangle) \xrightarrow{\ell} \text{Pss}(M_1 | \langle \text{route}(M_1); H' \rangle);$$

the internal satisfaction relation gives Σ'_1 s.t. $\Sigma_1 \xrightarrow{\ell} \Sigma'_1$ and $\models \text{Pss}(M_1 | \langle \text{route}(M_1); H' \rangle) : \Sigma'_1$. By coherence $\Sigma_2 \xrightarrow{\bar{\ell}} \Sigma'_2$ coherent with Σ'_1 . By definition of \mathcal{R} , we conclude.

- If $(M_1 | \langle \text{route}(M_1); H \rangle) \xrightarrow{\tau} N' = M'_1 | \langle \text{route}(M'_1); H' \rangle$ then we have

$$\text{Pss}(M_1 | \langle \text{route}(M_1); H \rangle) \xrightarrow{\tau} \text{Pss}(M_1 | \langle \text{route}(M_1); H \rangle).$$

By the internal satisfaction relation, we get Σ'_1 such that $\Sigma_1 \xrightarrow{\tau} \Sigma'_1$ and

$$\models \text{Pss}(M'_1 | \langle \text{route}(M'_1); H' \rangle) : \Sigma'_1.$$

By coherence $\Sigma_2 \xrightarrow{\tau} \Sigma'_2$ coherent with Σ'_1 . By definition of \mathcal{R} , we conclude.

Proposition 22 (Compositionality). If $\models M_i : \Sigma_i$ ($i = 1, 2$) such that Σ_1 and Σ_2 are coherent, then $\models M_1 | M_2 : \Sigma_1, \Sigma_2$.

Proof. We prove that $\{(M_1 | M_2, \Sigma_1, \Sigma_2)$ for all (M_1, M_2) verifying $\models M_i : \Sigma_i\}$ is a satisfaction relation.

1. Suppose $\Sigma_2, \Sigma_1 \xrightarrow{\ell} \Sigma'$. Exactly one of the two following statements is true:
 - $\Sigma_1 \xrightarrow{\ell} \Sigma'_1$. The satisfaction relation gives $M_1 \xrightarrow{\ell} M'_1$ such that $\models M'_1 : \Sigma'_1$, allowing us to conclude.
 - $\Sigma_2 \xrightarrow{\ell} \Sigma'_2$. We reason in a similar way.
2. Suppose $(M_1 | M_2) \xrightarrow{\ell} M'$. As $\ell \neq \tau$, exactly one of the two following statements is true:
 - $M_1 \xrightarrow{\ell} M'_1$. The satisfaction relation $\Sigma_1 \xrightarrow{\ell} \Sigma'_1$ and $\models M'_1 : \Sigma'_1$, allowing us to conclude.
 - $M_2 \xrightarrow{\ell} M'_2$. We reason in a similar way.
3. Suppose $(M_1 | M_2) \xrightarrow{\tau} M'$. Then either $M_1 \xrightarrow{\tau} M'_1$ or $M_2 \xrightarrow{\tau} M'_2$, and we reason as in the previous case, no interaction between two networks can take place, following the definition of the transition relation.

We define $n(\Sigma)$ with $\Sigma = (\alpha_i : \langle \Gamma_i; \Delta_i \rangle)_i$ to be the set of names present in at least one of the Δ_i .

Proposition 23 (Mixed compositionality). If $n(\Sigma_1) \cap n(\Sigma_2) = \emptyset$, $\models \Sigma_1, \Sigma_2 \triangleright N$ and $\models M : \Sigma_2$, then $\models \Sigma_1 \triangleright (\nu n(\Sigma_2))(N | M)$.

Proof. We prove that the relation

$$\mathcal{R} = \{(\nu n(\Sigma_2)) (N_1|M_2, \Sigma_1), \text{ for all } N_1, M_2, \Sigma_1 \text{ such that} \\ \models \Sigma_1, \Sigma_2 \triangleright N_1 \text{ and } \models M_2 : \Sigma_2\}$$

is a satisfaction relation.

- If $\Sigma_1 \xrightarrow{\ell} \Sigma'_1$ with ℓ output, then the external satisfaction relation gives $N_1 \xrightarrow{\bar{\ell}} N'_1$ and $\models \Sigma'_1, \Sigma_2 \triangleright N'_1$. As $n(\Sigma_1) \cap n(\Sigma_2) = \emptyset$, we have $(\nu n(\Sigma_2)) (N_1|M_2) \xrightarrow{\bar{\ell}} (\nu n(\Sigma_2)) (N'_1|M_2)$. The definition of \mathcal{R} allows us to conclude.
- If $(\nu n(\Sigma_2)) (N_1|M_2) \xrightarrow{\ell} N'$ with ℓ output, then it means that the subject of ℓ is not in $n(\Sigma_2)$. Suppose $M_2 \xrightarrow{\ell} M'_2$, this means $\Sigma_2 \xrightarrow{\ell} \Sigma'_2$: contradiction. Thus $N_1 \xrightarrow{\ell} N'_1$. The external transition relation gives $\Sigma_1 \xrightarrow{\bar{\ell}} \Sigma'_1$ and $\models \Sigma'_1, \Sigma_2 \triangleright N'_1$. The definition of \mathcal{R} allows us to conclude.
- If $(N_1|M_2) \xrightarrow{\tau} N'$, then
 - either $N_1 \xrightarrow{\tau} N'_1$ or $M_2 \xrightarrow{\tau} M'_2$, and we use the stability of the satisfaction relation by τ -transition to conclude;
 - or $N_1 \xrightarrow{\bar{\ell}} N'_1$ and $M_2 \xrightarrow{\ell} M'_2$.
 1. If ℓ is an output, then $M_2 \xrightarrow{\ell} M'_2$ gives us, by $M_2 : \Sigma_2$ and by the definition of partial satisfaction,

$$\Sigma_2 \xrightarrow{\ell} \Sigma'_2 \quad \text{such that} \quad \models M'_2 : \Sigma'_2 \quad (3)$$

We now combine the first part of (3) with $N_1 \xrightarrow{\bar{\ell}} N'_1$, as well as determinacy of specification transition (Prop. 20), to obtain, through $\models \Sigma_1, \Sigma_2 \triangleright N_1$,

$$\models \Sigma_1, \Sigma'_2 \triangleright N'_1 \quad (4)$$

By (3) and (4) as well as noting the name extrusion does not occur in the present system, the result is again in the closure, hence done.

2. If ℓ is an input, then we reason symmetrically: by $N_1 \xrightarrow{\bar{\ell}} N'_1$ (an output transition) and by the definition of full satisfaction, we get

$$\Sigma_2 \xrightarrow{\ell} \Sigma'_2, \quad \text{such that} \quad \models \Sigma_1, \Sigma'_2 \triangleright N'_1 \quad (5)$$

By the first part of (5), $M_2 \xrightarrow{\ell} M'_2$ and $\models M_2 : \Sigma_2$, we obtain

$$\models M'_2 : \Sigma'_2. \quad (6)$$

Again we note the result is in the closure again. We have exhausted all cases.