

Monitoring Networks through Multiparty Session Types

☆

Laura Bocchi^a, Tzu-chun Chen^b, Romain Demangeon^c, Kohei Honda^d, Nobuko Yoshida^e

^a*School of Computing, University of Kent, United Kingdom*

^b*Department of Computer Science, TU Darmstadt, Germany*

^c*LIP6, Université Pierre et Marie Curie, France*

^d*Department of Computer Science, Queen Mary, University of London, United Kingdom*

^e*Department of Computing, Imperial College London, United Kingdom*

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 recognised, 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.

Keywords: Session Types; the Pi - Calculus; Dynamic Monitoring; Runtime Verification; Bisimulation

☆This work has been partially sponsored by Ocean Observatories Initiative and EP-SRC EP/K011715/1, EP/G015635/1, EP/G015481/1, EP/K034413/1, EP/L00058X/1 and EP/N027833/1 and EU project FP7-612985 UpScale and COST Action IC1201 BETTY and ERC grant FP7-617805 LiVeSoft Lightweight Verification of Software.

1. Introduction

One of the main challenges in the engineering of 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, that was originally developed in the π -calculus [33, 6, 9, 23, 24, 16] 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 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. First, 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 type checking using session types. Unfortunately, it is not always straightforward to add these features to the specific host languages. Furthermore, the executable processes in a distributed system may be implemented in different languages. Second, for domains where dynamically typed or untyped languages are popular (e.g., Web programming), or in multi-organisational scenarios, the introduction of static typing infrastructure to support MPST may not be realistic.

Development of Heterogeneous Systems based on MPSTs. This article proposes a theoretical framework addressing the issues discussed above, by supporting the combination of static *and* dynamic verification of processes communicating in a network. Fig. 1 illustrates the proposed framework. As standard in MPST [33, 6], the first stage is to specify a global protocol as a *global type*, describing how the participants should interact in a multiparty session. The global type is then mechanically *projected* to generate local protocols, as *local types*, specifying the communication behaviour expected of each role in the session. The global type in Fig. 1 involves three roles, yielding three local types upon projection. Next, each principal in a network implements one (or possibly more) local types. We call these implementations endpoint processes, or simply processes. We aim to capture the decentralised nature of distributed application development, providing support for heterogeneous distributed systems by allowing components to be independently implemented, using different languages, libraries and programming techniques. Assume that (1) the process on the right-hand side of

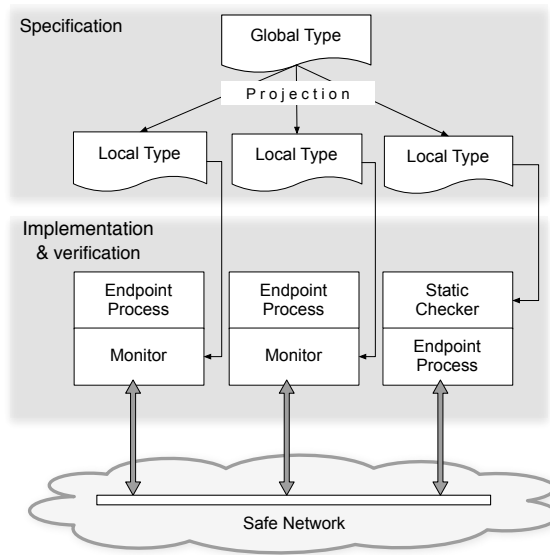


Figure 1: Static/dynamic verification through global types and projection

Fig. 1 is implemented in a language that supports static verification with session typing techniques, and that conformance to the implemented local type is verified this way, and (2) the other two processes are implemented in standard Java and Python, respectively, using simple session programming APIs and are not amenable to static typing. To ensure that the composition of these three processes conforms to the intended protocol we wrap the processes that cannot be statically verified with dedicated distributed *monitors*, that dynamically verify their participation in the session. In other words, our framework allows processes to be independently verified, either statically during deployment, or dynamically during execution, while retaining the strong global safety properties of statically verified systems.

This work is motivated in part by our ongoing collaboration with the Ocean Observatories Initiative (OOI) [44], 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. An implementation of the framework in Fig. 1 is currently integrated into the OOI infrastructure. In this implementation, processes are specified using Scribble [47, 51, 31, 30] (a practical incarnation of MPST) and processes are implemented in the Python programming language and dynamically monitored [43, 34].

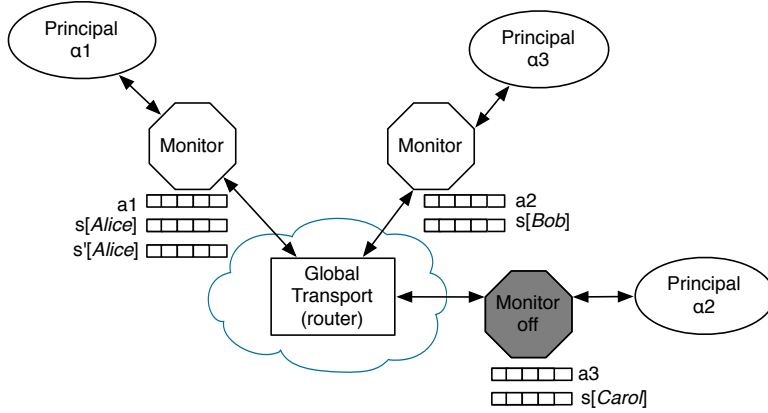


Figure 2: Architecture of monitored/unmonitored networks

Monitored Networks. Networks are organised as follows: a group of principals run processes communicating via asynchronous message passing; dedicated trusted monitors (one for each principal) guard the run-time behaviour of both the environment and that principal, through the evaluation of incoming and outgoing messages. The aim is to protect the principal from violations by other principals, and also to prevent the principal from committing violation (this can be used e.g. for debugging). Monitors regulate (1) the initiation, by principals, of new sessions, each specified by a well-defined global type, and (2) the movement of messages within each session. Fig. 2 illustrates the architecture of a network with three principals (α_1 , α_2 and α_3); all principals are monitored except α_3 , namely we assume the processes run by α_3 have been statically checked hence its monitor can be switched-off (indeed all outgoing and incoming messages can pass through without dynamic checking); each principal is associated with one or more *shared queues*, on which all other principals can send invitations to join new sessions. The messages exchanged within a session are all associated to *one common session ID*, and the exchange of messages in a session is regulated by verifying that the causality of messages follows the specification (roughly, the ensemble of local types) of that session. In Fig. 2 each principal is associated with exactly one shared queue, and we denote with a_i the queue associated with α_i , with $i \in \{1, 2, 3\}$; principal α_1 is currently playing role *Alice* in two sessions with session IDs s and s' , whereas α_2 and α_3 are playing *Bob* and *Carol*, respectively, in just one session s (e.g., the invitations to join s' have not yet been received by them).

A Formal Theory for Dynamic Verification. Our theory is based on the idea that, if the endpoint processes in a system are *independently* verified (either statically or dynamically) to conform a local type, then the corresponding global pro-

1
2
3
4
5 88 tocol is respected *as a whole*. To this goal, we propose a new formal model and
6 89 a bisimulation theory for heterogeneous networks of monitored and unmonitored
7 90 processes.

8 91 For the first time, we model dynamic verification based on types for the π -
9 92 calculus. We provide an explicit account of the *routing mechanism* that is im-
10 93 plicitly present inside the MPST framework: in a session, messages are sent to
11 94 abstract roles (e.g. to a Seller), and a router (a dynamically updated component of
12 95 the network) translates these roles into actual addresses.

13 96 Our approach also aims at giving a semantical equivalence for a collection of
14 97 protocols (and networks), by reaching a formal criterion for equating services. By
15 98 taking the routing feature into account when designing novel equivalences, our
16 99 formal model can relate networks built in different ways (through different distri-
17 100 butions or relocations of services) but offering the same *interface* to an external
18 101 observer. The router, being in charge of associating roles with principals, hides
19 102 to an external user the internal composition of a network: what distinguishes two
20 103 networks is not their structure but the services they are able to provide, or more
21 104 precisely, the local types they offer to the outside. We prove that bisimulation
22 105 is compositional (Proposition 4.4) and that equivalent networks satisfy the same
23 106 specification (Proposition 4.6).

24 107 We formally define a satisfaction relation to express when the behaviour of a
25 108 network conforms to a global specification and we prove a number of properties of
26 109 our model: *local safety* (Theorem 5.2) states that a monitored process respects its
27 110 local protocol, i.e. that dynamic verification by monitoring is sound; *global safety*
28 111 (Theorem 5.4) extends local safety to networks involving multiple principals; *lo-*
29 112 *cal transparency* (Theorem 6.1) states that a monitored process has equivalent
30 113 behaviour to an unmonitored but well-behaved (e.g. statically typed) process; and
31 114 *global transparency* (Theorem 6.3) states that a network where each principal is
32 115 monitored has equivalent behaviour to an unmonitored but well-behaved network.

33 116 Finally, we introduce a stronger property than global safety, *session fidelity*
34 117 (Theorem 7.13), which not only guarantees conformance of each monitored pro-
35 118 cess in a network to the ensemble of local specifications, but also requires that
36 119 the overall flow of messages throughout the router is correct. In this way, session
37 120 fidelity shows the correspondency between the behaviour of a monitored system
38 121 and the behaviour specified by a global protocol. Together, these properties justify
39 122 our framework for decentralised verification by allowing monitored and unmoni-
40 123 tored processes to be safely mixed while preserving protocol conformance for the
41 124 entire network. Technically, these properties also ensure the coherence of our the-
42 125 ory, by relating the satisfaction relations with the semantics and static validation
43 126 procedures.

44 127 Our theory is more involved than most of the existing works in the domain
45 128 of session verification [33] as, for the first time, both networks and monitoring

1
2
3
4
5 129 are made explicit. Our abstract model for session networks describe the evolution
6 130 of the network at a lower-level; for instance, we introduce dynamic update of
7 131 routing information: a participant taking part in a session does not send a message
8 132 to another participant, but sends a message to a role which is then routed by the
9 133 networks to the corresponding participant.

10
11
12 134 **Contributions and Outline.** This work is an extended version of [7] that in-
13 135 cludes: the definitions omitted in [7], additional examples, and full proofs. Specif-
14 136 ically, we extended [7] by including the following additional material:

- 17 137 • the formal definition of monitorability, a consistency condition on global
18 138 types, together with a discussion on its relevance and a statement of its
19 139 decidability (§ 2.2);
- 21 140 • the detailed definitions, full formal statement and proofs of session fidelity
22 141 and its relationship with global safety (in this introduction, § 5 and § 7),
23 142 which is only outlined in [7];
- 26 143 • a simpler but less restrictive semantics of networks (e.g., a principal is now
27 144 allowed to engage as different participants in the same session);
- 29 145 • a detailed formalisation for behavioural equivalences (§ 4.3)
- 31 146 • a formal statement on global safety in mixed (i.e., monitored *and* unmoni-
32 147 tored) networks (Corollary 6.4);

35 148 § 2 and § 3 introduce the formalisms for protocol specifications and networks, re-
36 149 spectively. § 3 provides a formal framework for monitored networks based on
37 150 π -calculus processes and protocol-based run-time enforcement through monitors.
38 151 § 4 introduces: a semantics for specifications (§ 4.1), a novel behavioural theory
39 152 for compositional reasoning over monitored networks through the use of equiva-
40 153 lences (bisimilarity and barbed congruence) and the satisfaction relation (§ 4.2).
41 154 Local and global safety are stated and proved in § 5, transparency in § 6, and ses-
42 155 sion fidelity in § 7. Related works are discussed in § 8 and future works in § 9.

47 156 2. Monitorability in Multiparty Session Types

48
49 157 This section provides basic definitions and well-formedness conditions for
50 158 multiparty session types. In § 2.1 we summarise the syntax of multiparty ses-
51 159 sion types annotated with logical assertions (MPST), which we use to model pro-
52 160 tocols. In § 2.2, we introduce a condition called *monitorability*, enforceable on
53 161 MPST, that sets the basis for the results presented in the next sections.

```

1
2
3
4
5
6   A ::= tt|ff|e1 = e2|e1 < e2|¬A|A1 ∧ A2|A1 ∨ A2
7   e ::= v | e1 + e2 | e1 - e2 | e1 * e2 | e1 mod e2
8
9   S ::= bool | int | string
10
11  G ::= r1 → r2 : {li(xi:Si){Ai}.Gi}i∈I | G1 | G2 | μt.G | t | end
12
13  T ::= r!{li(xi:Si){Ai}.Ti}i∈I | r?{li(xi:Si){Ai}.Ti}i∈I | μt.T | t | end
14

```

Figure 3: Global and local types with assertions

162 2.1. Multiparty Session Types with Assertions

163 Multiparty session types with assertions [9] are abstract descriptions of the
164 structure of interactions among the roles in a multiparty session (i.e., in a pro-
165 tocol); they specify the potential flows of messages, the conditions under which
166 interactions may occur, and the constraints on the communicated values.

167 Global types with assertions, or just *global types*, describe multiparty sessions
168 from a network perspective. Global types can be projected onto local types with
169 assertions, or just *local types*, each describing the protocol from the perspective
170 of a single role.

171 The syntax of global types (G, G', \dots) and local types (T, T', \dots) is defined
172 in Fig. 3. We let values v, v', \dots range over boolean constants, numerals and
173 strings, and e, e', \dots range over first-order expressions. *Assertions*, ranged over
174 by A, A', \dots are logical predicates used to express constraints on the values com-
175 municated. We consider assertions following the grammar given in Fig. 3 although
176 other decidable logics could be used. For instance, in [9, 24] the logics includes
177 existential quantifiers which we have omitted for simplicity (of evaluation of the
178 assertions by the run-time monitors), and because they are not necessary for our
179 run-time theory. The *sorts* of exchanged values (S, S', \dots) consists of atomic
180 types.

181 **Global Types with Assertions.** $r_1 \rightarrow r_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I}$ models an
182 interaction where role r_1 sends role r_2 one of the *branch labels* l_i , as well as a
183 payload denoted by an *interaction variable* x_i of sort S_i . Interaction variable x_i
184 binds its occurrences in A_i and G_i . A_i is the assertion which needs to hold for r_1
185 to select l_i , and which may constrain the values instantiating x_i . $G_1 \mid G_2$ specifies
186 two (independent) parallel threads in a session. We assume $G \mid \text{end}$ and $\text{end} \mid G$
187 are identical with G . $\mu t.G$ is a recursive type, where G is guarded in the standard
188 way [45, 6], and end ends the session.

54 *Example 2.1 (ATM: the global type).* Global type G_{ATM} specifies an ATM scenario.
55 Each session of ATM involves three roles: a client C, a payment server S and a
56

separate authenticator A.

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

189 At the beginning of the session C sends A payload x_i (i.e., the login details); then
190 A decides whether the authentication is successful or not, and informs S and C of
191 the choice by sending either label LoginOK or LoginFail. If LoginFail is cho-
192 sen then the session terminates. If LoginOK is chosen then C and S enter a loop
193 specified by G_{Loop} . In each iteration of G_{Loop} , S sends C the amount x_b currently
194 available in the account. The predicate states that x_b must be non negative. C can
195 then choose one of the following three labels: Withdraw (withdraws an amount
196 x_p , which must be positive and not exceed the current amount x_b), Deposit (de-
197 posits a positive amount x_d in the account), or Quit (ends the session). If either
198 Withdraw or Deposit was chosen then another iteration is executed.

199 **Local Types with Assertions.** Each local type T is associated with a role taking
200 part in a session. Local type $r!\{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I}$ models an interaction where
201 the role under consideration (say p) sends r a branch label l_i and a message de-
202 noted by an interaction variable x_i of sort S_i . Its dual is the receive interaction
203 $p?\{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I}$, where the role under consideration (say r) receives a
204 message from p. As customary for MPST, only global types can be composed in
205 parallel, namely there is no parallel composition of local types. This is guaranteed
206 by a well-formedness condition on global types (see Definition 2.4) formally de-
207 fined on the projection function, and requiring a given role to appear in only one
208 side of a parallel composition. The remaining local type syntax is similar to the
209 one of global types.

Example 2.2 (On causalities in global type). Consider the following global type:

$$\begin{aligned}
G_{\text{seq}} &= r_1 \rightarrow r_2 : (x : \text{int})\{\text{tt}\}. \\
&\quad r_3 \rightarrow r_4 : (y : \text{int})\{\text{tt}\}. \text{end}
\end{aligned}$$

The interaction from r_1 to r_2 and the interaction from r_3 to r_4 are causally un-
related. In fact, due to *distribution*, one cannot enforce r_3 to send y after r_2 has
received x (unless additional interactions are introduced, for example between r_2

and r_3). In fact (as in [33, 6, 9, 23, 24, 16]) the global type above specifies the same behaviour as

$$G_{par} = r_1 \rightarrow r_2 : (x : \text{int})\{\text{tt}\}.\text{end} \mid r_3 \rightarrow r_4 : (y : \text{int})\{\text{tt}\}.\text{end}$$

Where interactions are causally unrelated, we will use the parallel global types G_{par} rather than the sequential one G_{seq} .

The global type below, instead, requires the completion of the first interaction before r_2 can send the next message:

$$\begin{aligned} r_1 \rightarrow r_2 & : (x : \text{int})\{\text{tt}\}. \\ r_2 \rightarrow r_4 & : (y : \text{int})\{\text{tt}\}.\text{end} \end{aligned}$$

In addition, due to *asynchrony*, causality may affect the send and receive actions of an interaction in different ways, as shown by the global type below.

$$\begin{aligned} r_1 \rightarrow r_2 & : (x : \text{int})\{\text{tt}\}. \\ r_1 \rightarrow r_4 & : (y : \text{int})\{\text{tt}\}.\text{end} \end{aligned}$$

Variable y must be sent by r_1 only after variable x is *sent*, but possibly before x is *received* by r_2 .

One can derive a set of local types T_i from a global type G by *endpoint projection*. As in [23, 24], our definition of endpoint projection relies on a merge operator on local types which is useful to coherently assemble the behaviour that a role has in different branches of a global type, as illustrated in Example 2.3.

Example 2.3 (Merging local behaviours). Consider the following global type:

$$r_1 \rightarrow r_2 : \{ \begin{aligned} & l_1(x : \text{int})\{\text{tt}\}.r_2 \rightarrow r_3 : l_3(x' : \text{int})\{\text{tt}\}.\text{end}, \\ & l_2(y : \text{string})\{\text{tt}\}.r_2 \rightarrow r_3 : l_4(y' : \text{string})\{\text{tt}\}.\text{end} \end{aligned} \}$$

When defining the local behaviour of r_3 one must take into account that the first communication between r_1 and r_2 is not visible to r_3 . From the perspective of r_3 the session will either be described as *either* $r_2?l_3(x' : \text{int})\{\text{tt}\}.\text{end}$ or $r_2?l_4(y' : \text{string})\{\text{tt}\}.\text{end}$. The overall behaviour of r_3 is obtained by *merging* the two local types above into one branching as follows:

$$r_2?\{l_3(x' : \text{int})\{\text{tt}\}.\text{end}, l_4(y' : \text{string})\{\text{tt}\}.\text{end}\}$$

Definition 2.1. We define the union of two local types as the following partial operator:

1. $T \cup T = T$

$$2. \quad \mathbf{r}?l_k\{(x_k:S_k)\{A_k\}.T_k\}_{k \in I} \cup \mathbf{r}?l_k\{(x_k:S_k)\{A_k\}.T_k\}_{k \in J} = \\ \mathbf{r}?l_k\{(x_k:S_k)\{A_k\}.T_k\}_{k \in I \cup J} \quad \text{with } I \cap J = \emptyset$$

Local types are idempotent w.r.t. \cup and is otherwise undefined when types are not both input types from the same sender. In this case, all possible labels, together with their associated payload, types, assertions and continuations are collected into a single type.

Definition 2.2. Assume all labels are indexed and $l_i = l_j$ if and only if $i = j$. The merge operator \sqcup is a partial operator on local types, and is defined by the following axioms (closed by standard typed contexts):

$$1. \quad T \sqcup T = T \\ 2. \quad \mathbf{r}?l_i\{x_i:S_i\}\{A_i\}.T_i\}_{i \in I} \sqcup \mathbf{r}?l_j\{x'_j:S'_j\}\{A'_j\}.T'_j\}_{j \in J} = \\ \mathbf{r}?l_k\{(x_k:S_k)\{A_k\}.T_k\}_{k \in I \setminus J} \cup \mathbf{r}?l_k\{(x'_k:S'_k)\{A'_k\}.T'_k\}_{k \in J \setminus I} \cup \\ \mathbf{r}?l_k\{(x_k:S_k)\{A_k \vee A'_k\}.T_k \sqcup T'_k\}_{k \in I \cap J} \\ \text{when } \forall k \in I \cap J, x_k = x'_k, \text{ and } S_k = S'_k.$$

By (1) each local type is idempotent w.r.t. \sqcup . Axiom (2) merges two local types receiving messages from a common role, say \mathbf{r} . The resulting local type includes the union of the branches having distinguished labels (i.e. in $I \setminus J$ and $J \setminus I$), and integrates the common labels (i.e., in $I \cap J$). When integrating the common labels, axiom (2) makes sure that they have the same sorts (i.e., $S_k = S'_k$). The merge operator in [23, 24] is defined on local types without assertions. We define the predicate of the resulting local type to be the disjunction of the predicates of the local types being merged (i.e., $A_k \vee A'_k$). We motivate this choice via Examples 2.4 and 2.5. The intuition is that, when allowing merging for receiving actions only, the message, from the point of view of the local participant, satisfies a predicate for one of the branches.

Example 2.4 (Merging assertions). Consider the following global type:

$$\mathbf{r}_1 \rightarrow \mathbf{r}_2 \quad : \quad \{ \quad l_1(x : \text{int})\{\text{tt}\}.\mathbf{r}_2 \rightarrow \mathbf{r}_3 : l_3(x' : \text{int})\{x' > 0\}.\text{end}, \\ l_2(y : \text{string})\{\text{tt}\}.\mathbf{r}_2 \rightarrow \mathbf{r}_3 : l_3(x' : \text{int})\{x' > 10\}.\text{end} \}$$

This scenario differs from the one in Example 2.3 from the fact that \mathbf{r}_3 is expecting label l_3 in both branches (hence the behaviours of the common label l_3 need to be integrated). Role \mathbf{r}_3 must be able to accept a value for x' that satisfies $x' > 0$ or $x' > 10$ (without knowing which branch between l_1 or l_2 was selected, hence to which assertion \mathbf{r}_2 must obey). Therefore we relax the expectation of \mathbf{r}_3 to expect either case (i.e., a value satisfying the disjunction of the predicates):

$$\mathbf{r}_2?\{l_3(x' : \text{int})\{x' > 0 \vee x' > 10\}.\text{end}\}$$

1
2
3
4
5
243 When merging two local types, say T_1 and T_2 , if none of the axioms in Defini-
6
244 tion 2.2 applies then we say that T_1 and T_2 are non mergeable. As in [23, 24] we
7
245 let the local types for sending interactions to be non mergeable (i.e., axiom 2 can
8
246 only be applied to receive interactions) unless the send interactions to be merged
9
247 are identical (in which case one can merge by axiom 1). Example 2.5 illustrates
10
248 the motivation of this choice.

13
14 *Example 2.5 (Non mergeability of send interactions).* In the global type below

$$15 \quad \mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{ \begin{array}{l} 16 \quad l_1(x : \text{int})\{\mathbf{tt}\}.\mathbf{r}_3 \rightarrow \mathbf{r}_2 : l_3(x' : \text{int})\{\mathbf{tt}\}.\mathbf{end}, \\ 17 \quad l_2(y : \text{string})\{\mathbf{tt}\}.\mathbf{r}_3 \rightarrow \mathbf{r}_2 : l_4(y' : \text{string})\{\mathbf{tt}\}.\mathbf{end} \end{array} \}$$

19 Considering local behaviour of role \mathbf{r}_3 that role \mathbf{r}_3 must choose between l_3 and
20
21 l_4 without knowing which branch was chosen by \mathbf{r}_1 in the first interaction. If we
22
23 allowed to merge the two behaviours of \mathbf{r}_3 we would also allow, for instance, \mathbf{r}_3
24
25 to select l_3 after \mathbf{r}_1 had selected l_2 . In this scenario the behaviour \mathbf{r}_3 would not
26
27 conform to the expectations of \mathbf{r}_2 with respect to the global type. In Definition 2.2
28
29 we require, instead, that when a sender \mathbf{r} does not know which branch was chosen
30
31 by other roles in a previous interaction, then \mathbf{r} must act in the same way in all
32
33 branches. The following global type is, for instance, mergeable by axiom (1) in
34
35 Definition 2.2:

$$36 \quad \mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{ \begin{array}{l} 37 \quad l_1(x : \text{int})\{\mathbf{tt}\}.\mathbf{r}_3 \rightarrow \mathbf{r}_2 : l_3(x' : \text{int})\{\mathbf{tt}\}.\mathbf{end}, \\ 38 \quad l_2(y : \text{string})\{\mathbf{tt}\}.\mathbf{r}_3 \rightarrow \mathbf{r}_2 : l_3(x' : \text{int})\{\mathbf{tt}\}.\mathbf{end} \end{array} \}$$

39 Let $\text{roles}(G)$ be the set of roles in G . Formally,

$$40 \quad \begin{aligned} 41 \quad \text{roles}(\mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I}) &= \{\mathbf{r}_1, \mathbf{r}_2\} \cup_{i \in I} \text{roles}(G_i) \\ 42 \quad \text{roles}(G_1 \mid G_2) &= \text{roles}(G_1) \cup \text{roles}(G_2) \\ 43 \quad \text{roles}(\mu\mathbf{t}.G) &= \text{roles}(G) \\ 44 \quad \text{roles}(\mathbf{t}) = \text{roles}(\mathbf{end}) &= \emptyset \end{aligned}$$

45 We next define $\text{ftv}(G)$, the set of *free type variables* in G as:

$$46 \quad \begin{aligned} 47 \quad \text{ftv}(\mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I}) &= \cup_{i \in I} \text{ftv}(G_i) & \text{ftv}(\mathbf{end}) &= \emptyset \\ 48 \quad \text{ftv}(G_1 \mid G_2) &= \text{ftv}(G_1) \cup \text{ftv}(G_2) & \text{ftv}(\mu\mathbf{t}.G) &= \text{ftv}(G) \setminus \{\mathbf{t}\} & \text{ftv}(\mathbf{t}) &= \mathbf{t} \end{aligned}$$

49 The set $\text{fv}(A)$ of free variables occurring in A is defined as follows:

$$50 \quad \begin{aligned} 51 \quad \text{fv}(\mathbf{tt}) = \text{fv}(\mathbf{ff}) &= \emptyset & \text{fv}(e_1 = e_2) &= \text{fv}(e_1 < e_2) = \text{fv}(e_1) \cup \text{fv}(e_2) \\ 52 \quad \text{fv}(\neg A) &= \text{fv}(A) & \text{fv}(A_1 \wedge A_2) &= \text{fv}(A_1 \vee A_2) = \text{fv}(A_1) \cup \text{fv}(A_2) \\ 53 \quad \text{fv}(x) &= \{x\} \\ 54 \quad \text{fv}(e_1 \text{ op } e_2) &= \text{fv}(e_1) \cup \text{fv}(e_2) & \text{op} \in \{+, -, *, \text{mod}\} \end{aligned}$$

Definition 2.3 (Projection). Assume $r_1, r_2, r \in G$ and $r_1 \neq r_2$. The projection of G on r , written $G \upharpoonright r$, is defined as follows:

$$\begin{aligned}
(r_1 \rightarrow r_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I}) \upharpoonright r &= \\
&\begin{cases} r_2! \{l_i(x_i : S_i)\{A_i\}.(G_i \upharpoonright r)\}_{i \in I} & \text{if } r = r_1 \\ r_1? \{l_i(x_i : S_i)\{A_i\}.(G_i \upharpoonright r)\}_{i \in I} & \text{if } r = r_2 \\ \sqcup_{i \in I} G_i \upharpoonright r & \{r_1, r_2\} \cap \text{roles}(G_i) \neq \emptyset \\ G_i & G_i = \mathbf{t} \text{ or } G_i = \mathbf{end} \\ \text{undefined} & \text{otherwise} \end{cases} \\
(G_1 \mid G_2) \upharpoonright r &= \begin{cases} G_i \upharpoonright r & \text{if } \text{roles}(G_1) \cap \text{roles}(G_2) = \emptyset \\ & \text{and } \text{ftv}(G_1) \cap \text{ftv}(G_2) = \emptyset \\ & i \in \{1, 2\} \text{ and } r \notin \text{roles}(G_{3-i}) \\ \text{undefined} & \text{otherwise.} \end{cases} \\
\mu \mathbf{t}.G \upharpoonright r &= \begin{cases} \mu \mathbf{t}.(G \upharpoonright r) & \text{if } r \in G \\ \mathbf{end} & \text{otherwise} \end{cases} \\
\mathbf{t} \upharpoonright r &= \mathbf{t} \\
\mathbf{end} \upharpoonright r &= \mathbf{end}
\end{aligned}$$

249 The first rule projects an interaction onto sender or receiver role. Note that, if the
250 role is not involved in the interaction ($r \neq r_2 \neq r_1$) then the projection is the local
251 type resulting by merging (Definition 2.2) the projections $G_i \upharpoonright r$ for all $i \in I$. The
252 side condition ensures that we write a parallel composition if the roles and type
253 variables are disjoint (cf. Example 2.2). If some of the $G_i \upharpoonright r$ are non mergeable
254 then the projection rule cannot be applied. $(G_1 \mid G_2) \upharpoonright r$ is defined only when the
255 sets of roles of G_1 and G_2 are disjoint; in this case, the result of the projection is
256 the projection of the side of the parallel composition in which r appears (if r does
257 not appear further, then the projection of any side can only yield \mathbf{end}). The other
258 rules are straightforward.

259 If none of the rules in Definition 2.3 can be applied on a global type G then G
260 is *not projectable*.

261 **Definition 2.4 (Projectability).** A global type G is *projectable* if all of its projec-
262 tions to every role $r \in \text{roles}(G)$ are defined by Definition 2.3.

263 Hereafter in this article we will consider only projectable global types.

Example 2.6 (ATM: the local type of C). We present the *local type* T_C obtained by

projecting G_{ATM} on role C.

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

264 T_C specifies the behaviour that C should follow to meet the contract of global type
265 G_{ATM} . T_C states that C should first authenticate with A, then receive the message
266 Account from S, and then has the choice of sending Withdraw, or Deposit or
267 Quit. If label Withdraw or Deposit are chased another iteration is executed,
268 otherwise the session terminates.

269 2.2. Monitorability of Global Types

270 When designing a global type to be used in a monitoring framework, one
271 must ensure that the monitor associated to each role is always able to determine
272 if an incoming or outgoing message conforms to the contract or not. Example 2.7
273 shows that this is not the case for some global types.

Example 2.7 (Non monitorable global type). In the global type below r_3 does not know which value has been given to x in the first interaction between r_1 and r_2 .

$$\begin{array}{l}
G_s = r_1 \rightarrow r_2 : (x : \text{int})\{x > 5\}. \\
\quad r_2 \rightarrow r_3 : (y : \text{int})\{\text{tt}\}. \\
\quad r_3 \rightarrow r_1 : (z : \text{int})\{x > z\}.\text{end}
\end{array}$$

For any value sent by r_3 , the monitor of r_3 cannot determine whether the value sent for z by r_3 is violating or not. Similarly (but for receive interactions) in the global type below

$$\begin{array}{l}
G_r = r_1 \rightarrow r_2 : (x : \text{int})\{x > 5\}. \\
\quad r_2 \rightarrow r_4 : (y : \text{int})\{y > x\}.\text{end}
\end{array}$$

274 the monitor of r_4 will not have, at run-time, information on the value of variable
275 x , hence will not be able to determine if the value sent by r_2 for y conforms to the
276 assertion $y > x$.

277 We call global types as the ones illustrated in Example 2.7 *non monitorable*.
278 In the rest of this section we will give a formal definition of monitorability.

279 **Definition 2.5 (Known variables).** Let G'' be a subterm of G . We say that p
280 *knows* x in G'' if:

- 281 • there exists G' subterm of G s.t. $G' = r_1 \rightarrow r_2 : \{l_i(x_i : S_i)\{A_i\}.G_i\}_{i \in I}$,

- 282 • $p \in \{r_1, r_2\}$,
- 283 • for some $j \in I$, G'' is a subterm of G_j and $x = x_j$.

284 Namely, p knows a variable x in a subterm G'' of G if x is introduced in an
 285 interaction of G that occurs before G'' and that involves p .

286 **Definition 2.6 (Monitorability).** Let G be a subterm of G_0 . G is monitorable
 287 w.r.t. G_0 if one of the following conditions holds:

- 288 1. $G = r_1 \rightarrow r_2 : \{l_i(x_i : \text{int}_i)\{A_i\}.G_i\}_{i \in I}$, and for all $i \in I$, $y \in \text{fv}(A_i)$,
 289 $j \in \{1, 2\}$, r_j knows y in G and G_i is monitorable w.r.t. G_0 ;
- 290 2. $G = G_1 | G_2$, and for all $j \in \{1, 2\}$, G_j is monitorable w.r.t. G_0 ;
- 291 3. $G = \mu t. G'$ and G' is monitorable w.r.t. G_0 ;
- 292 4. $G = t$ or $G = \text{end}$.

293 We say that G is monitorable if it is monitorable w.r.t. G .

294 **Proposition 2.7 (Decidability).** Let G and G_0 be global types with G subterm of
 295 G_0 and $p \in \text{roles}(G_0)$. It is decidable if:

- 296 1. p knows x in G ,
- 297 2. G is monitorable w.r.t. G_0 .

298 *Proof.* (1) follows directly from: (i) the finiteness of the number of subterms of
 299 G_0 (and G), (ii) the finiteness of the number of labels in a branching (i.e., the
 300 cardinality of the set I of indices), and (iii) the decidability of inclusion in finite
 301 sets (e.g., $p \in \{r_1, r_2\}$ and $x \in \{x_i \mid i \in I\}$ in Definition 2.5). Proposition 2.7(2)
 302 follows from: (i) the finiteness of the number of subterms of G , (ii) the finiteness
 303 of the number of variables in assertions, and (iii) the decidability of p knows x in
 304 G by (1).

305 Knowledge of a name requires a linear search in the prefix. Monitorability
 306 requires a quadratic exploration.

307 In the following sections we will show that, under the assumption that all un-
 308 derlying global types are *projectable and monitorable*, the runtime monitoring
 309 discipline we propose ensures that the interactions in a session are safe (e.g., a
 310 principal implementing a role never receives messages of unexpected type), and
 311 predictable (i.e., faithful to the global interaction pattern specified by the proto-
 312 col).

313 Monitorability strengthens a similar property called history sensitivity in [9].
 314 By history sensitivity, only the sender of an interaction must know the free vari-
 315 ables of a predicate annotating that interaction. For instance, referring to Ex-
 316 ample 2.7, G_s is not history sensitive whereas G_r is. In [9] where only static

1
2
3
4
5
6 $P ::= \bar{a}\langle s[\mathbf{r}] : T \rangle \mid a(y[\mathbf{r}] : T).P \mid k[\mathbf{r}_1, \mathbf{r}_2]!l\langle e \rangle \mid k[\mathbf{r}_1, \mathbf{r}_2]?\{l_i(x_i).P_i\}_{i \in I} \mid$
7 $\text{if } e \text{ then } P \text{ else } Q \mid P \mid Q \mid \mathbf{0} \mid \mu X.P \mid X \mid P; Q \mid (\nu a) P \mid (\nu s)P$
8
9 $N ::= [P]_\alpha \mid N_1 \mid N_2 \mid \mathbf{0} \mid (\nu a)N \mid (\nu s)N \mid \langle r; h \rangle$
10
11 $r ::= \emptyset \mid r, s[\mathbf{r}] \mapsto \alpha \quad h ::= \emptyset \mid h \cdot m \quad m ::= \bar{a}\langle s[\mathbf{r}] : T \rangle \mid s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle$
12
13 $\mathbf{r}, \mathbf{r}_1, \dots$ roles s, s', \dots session names X, Y, \dots process variables
14 a, b, \dots shared names x, y, \dots variables P, Q, \dots processes
15 α, β, \dots principal names N, N', \dots networks
16
17
18
19
20

Figure 4: Processes and the network: syntax

21 317 verification is used, it is sufficient to check that all roles *send* values satisfying the
22 318 assertions. Receivers can rely on this fact thanks to the assumption that the pro-
23 319 cesses implementing the other roles are well-typed. In the run-time verification
24 320 scenario we cannot assume that the rest of the network behaves safely, hence both
25 321 sent and received values must be checked. The requirement posed by monitorabil-
26 322 ity, on the other hand, allows our theory to work using a logic without existential
27 323 quantifiers. On the contrary, in [9] quantifiers were needed, during endpoint pro-
28 324 jection, to close the assertions w.r.t. those variables that were unknown to the
29 325 receivers.
30
31
32
33

326 3. Formal Framework of Processes and Networks

327 In this section we introduce a novel *monitored session calculus* as a variant of
328 the π -calculus, which we use to model global networks. Global networks consists
329 of monitors and distributed programs, run by principals and implementing some
330 protocols.
31
32

331 In our formal framework, each distributed application consists of one or more
332 sessions among *principals*. A principal with behaviour P and name α is repre-
333 sented as $[P]_\alpha$. A *network* is a set of principals together with a (unique) *global*
334 *transport*, which abstractly represents the communication functionality of a dis-
335 tributed system. The syntax of processes, principals, and networks is given in Fig.
336 4, building on the multiparty session π -calculus from [6].
34
35

337 **Processes.** Processes, defined in Fig. 4, are ranged over by P, P', \dots and com-
338 municate using two types of channel: *shared channels* (or shared names) used
339 by processes for sending and receiving invitations to participate in sessions, and
340 *session channels* (or session names) used for communication *within* established
341 sessions. Each shared name, say a , is associated to one principal, say α , in the
342 sense that α can read from a ; a is shared in the sense that many other principals
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5 343 can send messages to α through a . One may consider shared names as e.g., URLs
6 344 or service names. The *session invitation* $\bar{a}\langle s[\mathbf{r}] : T \rangle$ invites, through a shared
7 345 name a , another process to play \mathbf{r} in a session s . The *session accept* $a(y[\mathbf{r}]:T).P$
8 346 receives a session invitation and, after instantiating y with the received session
9 347 name, behaves in its continuation P as specified by local type T for role \mathbf{r} . The
10 348 *selection* $k[\mathbf{r}_1, \mathbf{r}_2]!l\langle e \rangle$ sends, through session channel k (of an established ses-
11 349 sion), and as a sender \mathbf{r}_1 and to a receiver \mathbf{r}_2 , an expression e with label l . The
12 350 *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,
13 351 then behaves as P_i after instantiating x_i with the received value. We omit labels
14 352 when I is a singleton. The *conditional*, *parallel* and *inaction* are standard. The
15 353 *recursion* $\mu X.P$ defines X as P . Processes $(\nu a)P$ and $(\nu s)P$ hide shared names
16 354 and session names, respectively.

17
18
19
20
21 355 **Principals and Network.** Principals and networks are also formally defined in
22 356 Fig. 4. A *principal* $[P]_\alpha$, with process P and name α , represents a unit of be-
23 357 haviour (hence verification) in a distributed system. A *network* N is a collection
24 358 of principals with a unique global transport. The behaviour of a principal, de-
25 359 scribed in its process, includes communication over shared channels to create or
26 360 join new sessions, the communication over session channels, and control struc-
27 361 tures such as conditional branching and recursion.

30 362 A *global transport* is a pair $\langle r ; h \rangle$ of a routing table r that associates roles
31 363 to principals, and a global queue h . The *routing table* r is a finite map from
32 364 session-roles and shared names to principals. If, for instance, $r(a) = \alpha$ then a
33 365 session invitation message through a will be delivered to principal α . Similarly,
34 366 if $r(s[\mathbf{r}]) = \alpha$ then a message for \mathbf{r} in session s will be delivered to principal α .
35 367 The *global queue* h 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
36 368 over by m . These m represent messages-in-transit, i.e. those messages which have
37 369 been sent by some principal but have not yet been delivered. Possible shuffles
38 370 changing the ordering of in-transit messages is discussed below. Networks are
39 371 composed of principals and global transport.

42 372 Let n, n', \dots range over shared and session channels. A network N that sat-
43 373 isfies the following conditions is *well-formed*: (1) N contains at most one global
44 374 transport; (2) two principals in N never have the same principal name; and (3) if
45 375 $N \equiv (\nu \tilde{n})(\prod_i [P_i]_{\alpha_i} \langle r ; h \rangle)$ then each free shared or session name in P_i and h
46 376 occurs in \tilde{n} (we use $\prod_i P_i$ to denote $P_1 \mid P_2 \cdots \mid P_n$).

50 377 **Semantics.** The reduction relation for networks is generated from the rules de-
51 378 fined in Fig. 5, which model the interactions of principals with the global queue.
52 379 Rule $[\text{REQ}]$ places an invitation to participate as role \mathbf{r} in session s into the global
53 380 queue. Dually, in $[\text{ACC}]$, a process receives an invitation on a shared name from
54 381 the global queue, assuming a message on a is to be routed to α . As a result, the

$$\begin{array}{l}
1 \\
2 \\
3 \\
4 \\
5 \\
6 \quad [\bar{a}\langle s[\mathbf{r}] : T \rangle]_\alpha \mid \langle r ; h \rangle \longrightarrow [\mathbf{0}]_\alpha \mid \langle r ; h \cdot \bar{a}\langle s[\mathbf{r}] : T \rangle \rangle \quad [\text{REQ}] \\
7 \quad [a\langle y[\mathbf{r}] : T \rangle.P]_\alpha \mid \langle r ; \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot h \rangle \longrightarrow [P[s/y]]_\alpha \mid \langle r, s[\mathbf{r}] \mapsto \alpha ; h \rangle^\dagger \quad [\text{ACC}] \\
8 \quad [s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle]_\alpha \mid \langle r ; h \rangle \longrightarrow [\mathbf{0}]_\alpha \mid \langle r ; h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle \rangle \rangle \quad [\text{SEL}] \\
9 \\
10 \quad [s[\mathbf{r}_1, \mathbf{r}_2]?\{l_i(x_i).P_i\}_i]_\alpha \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 \langle r ; h \rangle^{\dagger\dagger} \quad [\text{BRA}] \\
11 \quad [\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}] \\
12 \\
13 \quad \frac{[P]_\alpha \mid N \longrightarrow [P']_\alpha \mid N'}{[\mathcal{E}(P)]_\alpha \mid N \longrightarrow [\mathcal{E}(P')]_\alpha \mid 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}] \\
14 \\
15 \\
16 \quad \dagger : r(a) = \alpha \quad \dagger\dagger : r(s[\mathbf{r}_2]) = \alpha \\
17 \\
18 \quad \mathcal{E} ::= () \mid \mathcal{E} \mid P \mid (\nu s)\mathcal{E} \mid (\nu a)\mathcal{E} \mid \mathcal{E}; P \mid \mathcal{E} \mid 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 \\
19 \\
20 \\
21 \\
22 \\
23 \\
24 \\
25 \\
26 \\
27 \\
28 \\
29 \\
30 \\
31 \\
32 \\
33 \\
34 \\
35 \\
36 \\
37 \\
38 \\
39 \\
40 \\
41 \\
42 \\
43 \\
44 \\
45 \\
46 \\
47 \\
48 \\
49 \\
50 \\
51 \\
52 \\
53 \\
54 \\
55 \\
56 \\
57 \\
58 \\
59 \\
60 \\
61 \\
62 \\
63 \\
64 \\
65
\end{array}$$

Figure 5: Reduction for dynamic networks

382 routing table adds $s[\mathbf{r}] \mapsto \alpha$ in the entry for s . Rule [SEL] puts in the queue a
383 message sent from \mathbf{r}_1 to \mathbf{r}_2 , which selects label l_j and carries v , if it is not going
384 to be routed to α (i.e. sent to self). Dually, [BRA] gets a message with label l_j from
385 the global queue, so that the j -th process P_j receives value v . Rules [CTX] are for
386 a closure under the reduction context \mathcal{E} . The other rules are standard.

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 [33, 23]. The rules are summarised in Fig. 6 where $u \in \{s, a\}$. The rule for message permutation is

$$\frac{m_1 \cdot m_2 \curvearrowright m_2 \cdot m_1}{h \cdot m_1 \cdot m_2 \cdot h' \equiv h \cdot m_2 \cdot m_1 \cdot h'}$$

387 and uses the notion of message permutation given in Definition 3.1. The rule for
388 message permutation is needed to treat the inherent non-determinism arising in
389 message orders when three or more participants are involved. The other rules are
390 straightforward.

391 **Definition 3.1 (Message permutation).** For messages in the global queue, we
392 say messages m_1 and m_2 are *permutable*, denoted by $m_1 \cdot m_2 \curvearrowright m_2 \cdot m_1$, if they
393 satisfy one of the following conditions:

- 394 1. $m_1 = s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle$ and $m_2 = s'\langle \mathbf{r}'_1, \mathbf{r}'_2, l'\langle v' \rangle \rangle$, where $s \neq s'$, or $(\mathbf{r}_1 \neq$
395 $\mathbf{r}'_1 \wedge \mathbf{r}_2 \neq \mathbf{r}'_2)$.
- 396 2. $m_1 = \bar{a}\langle s[\mathbf{r}] : T \rangle$ and $m_2 = \bar{a}'\langle s'[\mathbf{r}'] : T' \rangle$, where $a \neq a'$.
- 397 3. $m_i = \bar{a}\langle s[\mathbf{r}] : T \rangle$ and $m_j = s'\langle \mathbf{r}', \mathbf{r}'', l\langle v \rangle \rangle$ and $i, j \in \{1, 2\}, i \neq j$, where
398 $s \neq s'$, or $(\mathbf{r} \neq \mathbf{r}' \wedge \mathbf{r} \neq \mathbf{r}'')$.

$$\begin{array}{l}
1 \\
2 \\
3 \\
4 \\
5 \\
6 \quad P \mid \mathbf{0} \equiv P \quad P \mid Q \equiv Q \mid P \quad (P \mid Q) \mid R \equiv P \mid (Q \mid R) \\
7 \\
8 \\
9 \quad (\nu u)P \mid Q \equiv (\nu u)(P \mid Q) \text{ if } u \in \{s, a\}, u \notin \text{fn}(Q) \quad (\nu uu')P \equiv (\nu u'u)P \\
10 \\
11 \quad (\nu u)\mathbf{0} \equiv \mathbf{0} \quad \frac{P \equiv Q}{[P]_\alpha \equiv [Q]_\alpha} \quad (\nu u)[\mathbf{0}]_\alpha \equiv [\mathbf{0}]_\alpha \\
12 \\
13 \\
14 \quad N \mid [\mathbf{0}]_\alpha \equiv N \quad N_1 \mid N_2 \equiv N_2 \mid N_1 \quad (N_1 \mid N_2) \mid N_3 \equiv N_1 \mid (N_2 \mid N_3) \\
15 \\
16 \\
17 \quad (\nu u)N_1 \mid N_2 \equiv (\nu u)(N_1 \mid N_2) \text{ if } u \notin \text{fn}(N_2) \quad (\nu uu')N \equiv (\nu u'u)N \\
18 \\
19 \\
20 \quad \emptyset \cdot h \equiv h \quad \frac{m_1 \cdot m_2 \curvearrowright m_2 \cdot m_1}{h \cdot m_1 \cdot m_2 \cdot h' \equiv h \cdot m_2 \cdot m_1 \cdot h'} \quad \frac{r = r \quad h \equiv h'}{\langle r ; h \rangle \equiv \langle r' ; h' \rangle} \\
21 \\
22 \\
23 \\
24 \\
25
\end{array}$$

Figure 6: Structural congruence for networks

399 By (1) two interaction messages (i.e., in ongoing sessions) are permutable if they
400 are not related to the same session, or they are related to the same session but their
401 roles are different. By (2) two invitation messages are permutable if they are for
402 different principals. By (3) an invitation message and an interaction message are
403 permutable if they are for different sessions or the invited role r in the invitation
404 message is different from both the sender and receiver of the interaction message.

Example 3.1 (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 2.6) and communicating on session channel s .

$$\begin{array}{l}
39 \quad P_C = s[C, A]!\text{Login}(\text{"alice_pwd123"}); \\
40 \quad \quad s[A, C]?\{\text{LoginOK}(); \mu X.P'_C, \text{LoginFail}().\mathbf{0}\} \\
41 \quad P'_C = s[S, C]!\text{Account}(x_b); P''_C \\
42 \\
43 \quad \left. \begin{array}{l} \\ \\ \\ \end{array} \right| \begin{array}{l} P''_C = \text{if } \text{getmore}() \wedge (x_b \geq 10) \\ \quad \text{then } s[C, S]!\text{Withdraw}(10); X \\ \quad \text{else } s[C, S]!\text{Quit}(); \mathbf{0} \end{array}
\end{array}$$

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, and ff otherwise. P_S below implements the server role:

$$\begin{array}{l}
50 \\
51 \quad P_S = s[A, S]?\{\text{LoginOK}(); \mu X.P'_S, \text{LoginFail}().\mathbf{0}\} \\
52 \quad P'_S = s[S, C]!\text{Account}(\text{getBalance}()); P''_S \\
53 \\
54 \quad \left. \begin{array}{l} \\ \\ \end{array} \right| \begin{array}{l} P''_S = s[C, S]?\{\text{Withdraw}(x_p).X, \\ \quad \text{Deposit}(x_d).X, \\ \quad \text{Quit}().\mathbf{0}\} \end{array}
\end{array}$$

405 We assume that $\text{getBalance}()$ is a function, local to the principal running P_S , that
406 synchronously returns the current balance of the client.

407 4. The Monitored Network: Semantics and Equivalences

408 In this section we formalise the specifications (based on local types) used to
 409 guard the runtime behaviour of the principals in a network. These specifications
 410 are the foundation of system monitors, each wrapping a principal to ensure that
 411 the ongoing communication conforms to the given specification. Then, we present
 412 a behavioural theory for monitored networks and their safety properties.

413 4.1. Semantics of Global Specifications

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

$$418 \Gamma ::= \emptyset \mid \Gamma, a : \mathbb{I}(T[\mathbf{r}]) \mid \Gamma, a : \mathbb{O}(T[\mathbf{r}]) \quad \Delta ::= \emptyset \mid \Delta, s[\mathbf{r}] : T$$

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

$$425 \Sigma ::= \emptyset \mid \Sigma, \alpha : \langle \Gamma; \Delta \rangle$$

426 The semantics of Σ is defined using the following labels:

$$427 \ell ::= \bar{a}\langle s[\mathbf{r}] : T \rangle \mid a\langle s[\mathbf{r}] : T \rangle \mid s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle \mid s[\mathbf{r}_1, \mathbf{r}_2]?l\langle v \rangle \mid \tau$$

428 The first two labels are for *invitation* actions, the first is for requesting and the sec-
 429 ond is for accepting. Labels with $s[\mathbf{r}_1, \mathbf{r}_2]$ indicate *interaction* actions for sending
 430 (!) or receiving (?) messages within sessions. The labelled transition relation for
 431 specifications is defined by the rules in Fig. 7 Rule [REQ] allows α to send an invi-
 432 tation on a properly typed shared channel a (i.e., given that the shared environment
 433 maps a to $T[\mathbf{r}]$). Rule [ACC] allows α to receive an invitation to be role \mathbf{r} in a new
 434 session s , on a properly typed shared channel a . Rule [BRA] allows α , participating
 435 to session s as \mathbf{r}_2 , to receive a message with label l_j from \mathbf{r}_1 , given that A_j is satis-
 436 fied after replacing x_j with the received value v . After the application of this rule
 437 the specification is T_j . Rule [SEL] is the symmetric (output) counterpart of [BRA].
 438 We use \downarrow to denote the evaluation of a logical assertion. [SPL] is the juxtaposition
 439 of two session environments. [TAU] says that the specification should be invariant
 440 under reduction of principals. [PAR] says if Σ_1 and Σ_3 are composable, after Σ_1
 441 becomes as Σ_2 , they are still composable.

$$\begin{array}{c}
1 \\
2 \\
3 \\
4 \\
5 \\
6 \quad \alpha : \langle \Gamma, a : 0(T[\mathbf{r}]); \Delta \rangle \xrightarrow{\bar{a}\langle s[\mathbf{r}]:T \rangle} \alpha : \langle \Gamma, a : 0(T[\mathbf{r}]); \Delta \rangle \quad [\text{REQ}] \\
7 \\
8 \quad \frac{s \notin \text{dom}(\Delta)}{\alpha : \langle \Gamma, a : \text{I}(T[\mathbf{r}]); \Delta \rangle \xrightarrow{a\langle s[\mathbf{r}]:T \rangle} \alpha : \langle \Gamma, a : \text{I}(T[\mathbf{r}]); \Delta, s[\mathbf{r}]:T \rangle} \quad [\text{ACC}] \\
9 \\
10 \quad \frac{\Gamma \vdash v : S_j, A_j[v/x_j] \downarrow \text{tt}, j \in I}{\alpha : \langle \Gamma; \Delta, s[\mathbf{r}_2] : \mathbf{r}_1 ? \{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I} \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2] ? l_j \langle v \rangle} \alpha : \langle \Gamma; \Delta, s[\mathbf{r}_2] : T_j[v/x_j] \rangle} \quad [\text{BRA}] \\
11 \\
12 \quad \frac{\Gamma \vdash v : S_j, A_j[v/x_j] \downarrow \text{tt}, j \in I}{\alpha : \langle \Gamma; \Delta, s[\mathbf{r}_1] : \mathbf{r}_2 ! \{l_i(x_i : S_i)\{A_i\}.T_i\}_{i \in I} \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2] ! l_j \langle v \rangle} \alpha : \langle \Gamma; \Delta, s[\mathbf{r}_1] : T_j[v/x_j] \rangle} \quad [\text{SEL}] \\
13 \\
14 \quad \frac{\alpha : \langle \Gamma_1; \Delta_1 \rangle \xrightarrow{\ell} \alpha : \langle \Gamma'_1; \Delta'_1 \rangle}{\alpha : \langle \Gamma_1; \Delta_1, \Delta_2 \rangle \xrightarrow{\ell} \alpha : \langle \Gamma'_1; \Delta'_1, \Delta_2 \rangle} \quad \Sigma \xrightarrow{\tau} \Sigma \quad \frac{\Sigma_1 \xrightarrow{\ell} \Sigma_2}{\Sigma_1, \Sigma_3 \xrightarrow{\ell} \Sigma_2, \Sigma_3} \quad [\text{SPL,TAU,PAR}] \\
15 \\
16 \\
17 \\
18 \\
19 \\
20 \\
21 \\
22 \\
23 \\
24 \\
25 \\
26 \\
27 \\
28 \\
29 \\
30 \\
31 \\
32 \\
33 \\
34 \\
35 \\
36 \\
37 \\
38 \\
39 \\
40 \\
41 \\
42 \\
43 \\
44 \\
45 \\
46 \\
47 \\
48 \\
49 \\
50 \\
51 \\
52 \\
53 \\
54 \\
55 \\
56 \\
57 \\
58 \\
59 \\
60 \\
61 \\
62 \\
63 \\
64 \\
65
\end{array}$$

Figure 7: Labelled transition relation for specifications

4.2. 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:

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

The reduction rules for monitored networks are given in Fig. 8 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 § 4.1).

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, [BRA] (resp. [SEL]) extracts (resp. introduces) messages from (resp. in) the global queue. The error cases for [REQ] and [SEL], namely [REQER] and [SELER], ‘skip’ the current action (removing it from the process), and do not modify the queue, the router nor the state of the monitor. The error cases for [ACC] and [BRA], namely [ACCER] and [BRAER], do not affect the process, which remains ready to perform the action, and remove the violating message from the queue.

Example 4.1 (ATM: a monitored network). We illustrate the monitored network for the ATM scenario, where the routing table is defined as

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

$$\begin{array}{c}
1 \\
2 \\
3 \\
4 \\
5 \\
6 \\
7 \\
8 \\
9 \\
10 \\
11 \\
12 \\
13 \\
14 \\
15 \\
16 \\
17 \\
18 \\
19 \\
20 \\
21 \\
22 \\
23 \\
24 \\
25 \\
26 \\
27 \\
28 \\
29 \\
30 \\
31 \\
32 \\
33 \\
34 \\
35 \\
36 \\
37 \\
38 \\
39 \\
40 \\
41 \\
42 \\
43 \\
44 \\
45 \\
46 \\
47 \\
48 \\
49 \\
50 \\
51 \\
52 \\
53 \\
54 \\
55 \\
56 \\
57 \\
58 \\
59 \\
60 \\
61 \\
62 \\
63 \\
64 \\
65
\end{array}$$

$$\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\langle y[\mathbf{r}] : T \rangle . 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\langle y[\mathbf{r}] : T \rangle . P]_\alpha \mid M \mid \langle r ; \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot h \rangle \longrightarrow [a\langle y[\mathbf{r}] : T \rangle . P]_\alpha \mid M \mid \langle r ; h \rangle} \\
\text{[BRAER]} \frac{M \xrightarrow{\ell} \ell = s[\mathbf{r}_1, \mathbf{r}_2]?l_j\langle v \rangle}{[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 [s[\mathbf{r}_1, \mathbf{r}_2]? \{l_i(x_i).P_i\}_i]_\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}$$

Figure 8: Reduction for monitored networks (assume $M = \alpha : \langle \Gamma ; \Delta \rangle$).

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 3.1, and the process of A is $\mathbf{0}$.

$$\begin{array}{l}
N_S = [P'_S]_\alpha \mid M_S = [s[S, C]! \text{Account}\langle 100 \rangle; P'_S]_\alpha \mid M_S \text{ (assuming } \textit{getBalance}() = 100) \\
N_C = [P'_C]_\beta \mid M_C = [s[S, C]? \text{Account}(x_b).P'_C]_\beta \mid M_C \\
N_A = [\mathbf{0}]_\gamma \mid \gamma : \langle c : T_A[A] ; s[A] : \text{end} \rangle
\end{array}$$

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

$$\begin{array}{l}
N_1 = [s[S, C]! \text{Account}\langle 100 \rangle; P'_S]_\alpha \mid M_S \mid [s[S, C]? \text{Account}(x_b).P'_C]_\beta \mid M_C \mid N_A \mid \langle r ; \emptyset \rangle \\
\longrightarrow [P'_S]_\alpha \mid M'_S \mid [P'_C[100/x_b]]_\beta \mid M'_C \mid N_A \mid \langle r ; \emptyset \rangle
\end{array}$$

449 where $M'_S = \alpha : \langle a : T_S[S] ; s[S] : T'_S \rangle$ and $M'_C = \beta : \langle b : T_C[C] ; s[C] : T'_C \rangle$

450 Above, $x_b \geq 0$ is satisfied since $x_b = 100$. If the server tried to communicate e.g.,
451 value -100 for x_b , then the monitor (by rule [SELER]) would drop the message.

452 Following Example 4.1, in the example that follows we show the different
453 behaviours of monitored and unmonitored processes.

Example 4.2 (Compare a monitored process to an unmonitored one.). Let

$$\begin{array}{l}
\ell = s[S, C]! \text{Account}\langle -10 \rangle \\
P_1 = s[S, C]! \text{Account}\langle -10 \rangle; P'_S \\
M_S = \alpha : \langle a : T_S[S] ; s[S] : C! \text{Account}(x_b : \text{int})\{x_b \geq 0\}.T'_S \rangle
\end{array}$$

1
2
3
4
5 454 The unmonitored principal $[P_1]_\alpha$ can make a step ℓ , namely $[P_1]_\alpha \xrightarrow{\ell}$. However,
6 455 the its monitored counter-part $N_S = [P_1]_\alpha \mid M_S$ cannot make a step ℓ (that is
7
8 456 $N_S \not\xrightarrow{\ell}$) since the value -10 does not satisfy the predicate $x_b \geq 0$ attached to the
9 457 local type of the session monitored by M_S .

Similarly, for type violations, consider:

$$\begin{aligned} \ell &= s[S, C]! \text{Account} \langle \text{“hello”} \rangle \\ P_2 &= s[S, C]! \text{Account} \langle \text{“hello”} \rangle; P''_S \end{aligned}$$

12
13
14
15
16 458 then $[P_2]_\alpha \xrightarrow{\ell}$ but $[P_2]_\alpha \mid M_S \not\xrightarrow{\ell}$.

18 459 4.3. Network Satisfaction and Equivalences

20 460 Based on the formal representations of monitored networks, we now intro-
21 461 duce the key formal tools for analysing their behaviour. Concretely, we introduce
22 462 two different equivalences to semantically compare networks: *bisimulation* and
23 463 *barbed congruence* (the latter relying on the notion of *interface*, also given in
24 464 this section). The two equivalences allow us to compare networks using different
25 465 granularities. On the one side, bisimilarity addresses mainly partial networks, and
26 466 gives an equivalence that distinguishes two networks containing different com-
27 467 ponents. On the other side, (barbed) congruence addresses networks (includ-
28 468 ing global transport) from the point of view of an external observer; thus, two
29 469 networks built from different components but offering the same service will be
30 470 equated. We choose to give two equivalences in order to give the theory a way to
31 471 compare session networks from two different points of views: bisimulation allows
32 472 designers to equate networks whose structures are similar, whereas barbed con-
33 473 gruence allows users to equate networks, seen as black boxes, which provide the
34 474 same service. Both equivalences are compositional, as proved in Proposition 4.4.
35 475 Finally, using the definition of congruence, we define *the satisfaction relation*
36 476 $\models N \triangleright M$, used in §5 and §7 to prove the properties of our framework.

37 477 **Bisimulations.** We first define semantics for networks of components, or partial
38 478 networks, on which we define bisimulation: we use M, M', \dots for a *partial net-*
39 479 *work*, that is a network without a global transport, hence allowing the global ob-
40 480 *serva-*tion of interactions. The labelled transition relation for processes and partial
41 481 networks M is defined in Fig. 9. In (CTX), $n(\ell)$ indicates the names occurring in
42 482 ℓ while $\text{bn}(\mathcal{E})$ indicates binding induced by \mathcal{E} . In (RES), $\text{sbj}(\ell)$ denotes the subject
43 483 of ℓ . In (TAU) the axiom is obtained either from the reduction rules for dynamic
44 484 networks given in § 3 (only those not involving the global transport), or from the
45 485 corresponding rules for monitored networks (which have been omitted in § 4.2).

46 486 Hereafter we write \Longrightarrow for $\xrightarrow{\tau^*}$, $\xRightarrow{\ell}$ for $\xrightarrow{\ell} \Longrightarrow$, and $\xRightarrow{\hat{\ell}}$ for $\xrightarrow{\ell} \xRightarrow{\ell}$
47 487 and $\xRightarrow{\ell}$ otherwise.

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

Figure 9: Labelled transition relation for processes and partial networks

Definition 4.1 (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 $\mathbf{bn}(\ell) \cap \mathbf{fn}(M_2) = \emptyset$, 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. As stated above, we build another model where two different implementations of the same service are equated. Bisimilarity is too strong for this aim as shown in further Example 4.3. We therefore introduce a contextual congruence (barbed reduction-closed congruence [32]) \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 that the two (barbed-congruent) networks propose the same *interface* to the exterior. More precisely, two networks are related with \cong when, composed with the same third network, they offer the same *barbs* (i.e., 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 4.2 (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 $N \longrightarrow^* N' \downarrow_a$.

We denote $\mathcal{P}(N)$ for the set of principals in N , that is $\mathcal{P}(\prod_i [P_i]_{\alpha_i}) = \{\alpha_1, \dots, \alpha_n\}$. 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 $N_i = (\nu \tilde{n}_i)(M_i \mid \langle r_i ; h_i \rangle)$ ($i = 1, 2$).

Definition 4.3 (Barbed reduction-closed congruence). A relation \mathcal{R} on networks with the same principals is a *barbed r.c. congruence* [32] if the following holds:

1
2
3
4
5 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)
6 $N_1 \longrightarrow N'_1$ implies $N_2 \longrightarrow^* N'_2$ s.t. $N'_1 \mathcal{R} N'_2$ again, and the symmetric case; (3)
7 $N'_1 \Downarrow_a$ iff $N'_2 \Downarrow_a$. We write $N_1 \cong N_2$ when N_1 and N_2 are related by a barbed r.c.
8 congruence.
9

10
11 **Properties.** The following result states that composing two bisimilar partial net-
12 works with the same network – implying the same router and global transport –
13 yields two undistinguishable networks.
14

15
16 **Proposition 4.4 (Congruency).** If $M_1 \approx M_2$, then (1) $M_1|M \approx M_2|M$ for each
17 composable partial network M ; and (2) $M_1|N \cong M_2|N$ for each composable
18 network N .
19

20 *Proof.* For (1) we show that the relation

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

21
22
23
24
25 is a bisimulation. Suppose $(M_1|M) \mathcal{R} (M_2|M)$ and $M_1|M \xrightarrow{\ell} M^1$. We discuss the
26 shape of M^1 :
27

- 28 • If $M^1 = M'_1|M$, it means that $M_1 \xrightarrow{\ell} M'_1$. By definition of \mathcal{R} , $M_2 \xrightarrow{\hat{\ell}} M'_2$
29 and $M'_1 \approx M'_2$, we conclude.
30
31
- 32 • If $M^1 = M_1|M'$, it means that $M \xrightarrow{\ell} M'$. It is easy to conclude.
33

34
35 By examining the reduction rule associated to parallel composition, we observe
36 that no reduction is induced through interactions between the two networks. Hence
37 we have covered all cases. The symmetric case (when $M_2|M \xrightarrow{\ell} M^2$) is similar.
38

39 To prove (2) we proceed by showing that

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

40
41
42
43 is a barbed congruence. First, \mathcal{R} is clearly a congruence since it is closed under
44 composition. Second, for (2), we take a composable N' . We have $N' \diamond (M_i|N) =$
45 $M_i|(N' \diamond N)$ for $i \in \{1, 2\}$. We use the definition of \mathcal{R} to conclude. For (3),
46 assume $M_1|N \longrightarrow N^1$.
47

- 48
49 • If $N^1 = M_1|N'$, it means that $N \longrightarrow N'$. We use the definition of \mathcal{R} to
50 conclude.
51
- 52 • If $N^1 = M'_1|N'$, it means that $N = M_0|\langle r ; \ell \cdot H \rangle$, $N = M'_0|\langle r ; H \rangle$
53 and $M_1 \xrightarrow{\ell} M'_1$. We deduce $N^2 = M'_2|N'$, with $N = M_0|\langle r ; \ell \cdot H \rangle$,
54 $N = M'_0|\langle r ; H \rangle$ and $M_2 \xrightarrow{\ell} M'_2$. We use the definition of \mathcal{R} to conclude.
55
56

541 • If the reduction is induced by interaction between M_1 and N , then M_2 has
 542 the corresponding action, hence we can reason in the same way, hence done.

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

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

546 The symmetric case is similar.

547 By definition this shows $\approx \subset \cong$.

548 *Example 4.3 (Example of Equivalence)*. We give here an example illustrating our
 549 equivalences of networks. Consider the following networks:

$$\begin{aligned}
 M'_0 &= [a_1(y_1[C] : T_C).y_1[C, A]!\langle \text{Login}(x_i) \rangle.y_1[A, C]?\text{LoginOK}().P_{\text{LOOP},C}]_{\alpha_1} \\
 M'_1 &= [a_2(y_2[S] : T_S).P_{\text{LOOP},S}]_{\alpha_2} \\
 &\quad | [(\nu s) \bar{a}_1\langle s[C] : T_C \rangle \mid \bar{a}_2\langle s[S] : T_S \rangle \mid \bar{a}_3\langle s[A] : T_A \rangle \\
 &\quad \mid a_3(y_3[A] : T_A).y_3[C, A]?\langle \text{Login}(x_i) \rangle.y_3[A, S]!\langle \text{LoginOk}() \rangle].P_{\text{LOOP},A}]_{\beta} \\
 M'_2 &= [a_2(y_2[S] : T_S).P_{\text{LOOP},S}]_{\alpha_2} \\
 &\quad | [a_4(y_4[\text{DB}] : T_{\text{DB}}).y_4[A, \text{DB}]?\langle \text{Query} \rangle.y_4[\text{DB}, A]!\langle \text{Answer} \rangle]_{\gamma} \\
 &\quad | [(\nu s) (\bar{a}_1\langle s[C] : T_C \rangle \mid \bar{a}_2\langle s[S] : T_S \rangle \mid \bar{a}_3\langle s[A] : T_A \rangle \mid \bar{a}_4\langle s[\text{DB}] : T_{\text{DB}} \rangle) \\
 &\quad \mid a_3(y_3[A] : T'_A).y_3[C, A]?\langle \text{Login}(x_i) \rangle. \\
 &\quad \quad y_3[A, \text{DB}]!\langle \text{Query} \rangle.y_3[\text{DB}, A]?\langle \text{Answer} \rangle].P_{\text{LOOP},A}]_{\beta} \\
 N'_1 &= M'_1 \mid \langle a_1 \mapsto \alpha_1, a_2 \mapsto \alpha_2, a_3 \mapsto \beta ; \emptyset \rangle \\
 N'_2 &= (\nu a_4) (M'_2 \mid \langle a_1 \mapsto \alpha_1, a_2 \mapsto \alpha_2, a_3 \mapsto \beta, a_4 \mapsto \gamma ; \emptyset \rangle)
 \end{aligned}$$

550 Our networks implement the ATM example defined in 2.1. For the sake of
 551 clarity, we have to take the following shortcuts: (1) we only consider the login
 552 phase of the protocol, the LOOP phase is abstracted into three processes $P_{\text{LOOP},C}$,
 553 $P_{\text{LOOP},A}$, $P_{\text{LOOP},S}$ for the three different roles, (2) to lighten the notations, we do not
 554 make the logical annotations explicit, (3) as a result of (2), we do not implement
 555 login validation and only write the case were the login succeeds.

556 We present two different networks N'_1 and N'_2 , both are implementing the
 557 Server-Authenticator part of the ATM protocol. The Server part is the same in
 558 both processes (executed at principal α_2), but the Authenticator part (executed at
 559 β) is different: N'_1 implements straightforwardly the protocol while N'_2 contains
 560 another indirection involving a fourth participant (executed at γ): the Authentica-
 561 tor sends a query to a *Database* to retrieve additional information required in the
 562 login process, and the Database answers.

563 Thus, the protocols implemented in both networks are different, as one in-
 564 volves three participants and the other one four. Yet, the query to the Database

1
2
3
4
5 is N'_2 is unobservable from the outside, and an external client, such as N'_0 cannot
6 distinguish between N'_1 and N'_2 .

7 This is captured by our equivalences: the two partial networks M'_1 and M'_2 do
8 not contain the same components, and as a result, are not bisimilar: after some
9 steps, M'_2 is able to emit on the channel a_4 , which is impossible for M'_1 . How-
10 ever, when encapsulated into dynamic networks N'_1 and N'_2 , they are barbed r.c.
11 congruent: they will offer it the same interface to the same external client.
12
13

14 **Satisfaction.** We finally present a satisfaction relation for partial networks that
15 include local principals. If M is a partial network, $\models M \triangleright \Sigma$ s.t. $\text{dom}(\Sigma) = \mathcal{P}(M)$,
16 it means that: the specification Σ allows all the outputs from the network; the
17 network M is ready to receive all the inputs indicated by the specification; and
18 this is preserved by transition.
19
20

21 **Definition 4.5 (Satisfaction).** Let $\text{subj}(\ell)$ denote the subject of $\ell \neq \tau$. A relation
22 \mathcal{R} from partial networks to specifications is a *satisfaction* when $M\mathcal{R}\Sigma$ implies:
23
24

- 25 1. If $\Sigma \xrightarrow{\ell} \Sigma'$ for an input ℓ and M has an input at $\text{subj}(\ell)$, then $M \xrightarrow{\ell} M'$ s.t.
26 $M'\mathcal{R}\Sigma'$.
27
28 2. If $M \xrightarrow{\ell} M'$ for an output at ℓ , then $\Sigma \xrightarrow{\ell} \Sigma'$ s.t. $M'\mathcal{R}\Sigma'$.
29
30 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).
31

32 When $M\mathcal{R}\Sigma$ for a satisfaction relation \mathcal{R} , we say M *satisfies* Σ , denoted $\models M \triangleright$
33 Σ . By Definition 4.5 and Proposition 4.4 we obtain:
34

35 **Proposition 4.6 (Satisfaction).** If $M_1 \cong M_2$ and $\models M_1 \triangleright \Sigma$ then $\models M_2 \triangleright \Sigma$.

36 That is, if two networks present the same interface, they satisfy the same specifi-
37 cations.
38
39

40 5. Safety Assurance in Partial Networks

41
42 In this section, we present the properties underpinning safety assurance in
43 partial networks, that are networks without a global transport. By considering
44 partial networks we focus on the properties of principals (and their respective
45 monitors) with respect to specifications, abstracting from the routing mechanisms.
46 The routing mechanisms will be taken into account in later sections. We first
47 consider networks consisting of single monitored principals (*local safety*) and then
48 extend the results to partial networks in general (*global safety*).
49
50

51 Recall that: partial networks are networks without global transport; M denotes
52 an unmonitored partial network; N denotes an unmonitored network; \mathbb{N} denotes
53 a (monitored or unmonitored) network. Monitors, ranged over by M , are specifi-
54 cations (of the form $\alpha : \langle \Gamma; \Delta \rangle$) used for dynamic verification. See Table 5 for a
55 summary of the notation.
56
57
58
59
60
61
62
63
64
65

M	monitor	specification used for monitoring
M	unmonitored partial network	without M, without $\langle r ; h \rangle$
N	unmonitored network	without M, with $\langle r ; h \rangle$
N	network	with or without M, with $\langle r ; h \rangle$

Table 1: Networks: summary of the notations

The partial network composed by a principal guarded by its monitor can take any action expected by the specification:

Lemma 5.1. For any principal $[P]_\alpha$, specification $\alpha : \langle \Gamma, \Delta \rangle$, and action ℓ , if $\alpha : \langle \Gamma, \Delta \rangle \xrightarrow{\ell} \alpha : \langle \Gamma', \Delta' \rangle$ and $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$, then $[P]_\alpha \mid \alpha : \langle \Gamma, \Delta \rangle \xrightarrow{\ell} [P']_\alpha \mid \alpha : \langle \Gamma', \Delta' \rangle$.

Proof. Direct, as no interaction can appear between $[P]_\alpha$ and its monitor with specification $\alpha : \langle \Gamma, \Delta \rangle$ when ℓ is performed. \square

Local safety ensures that a monitored process always behaves well with respect to the specification used to define its monitor.

Theorem 5.2 (Local safety). $\models [P]_\alpha \mid M \triangleright \alpha : \langle \Gamma; \Delta \rangle$ with $M = \alpha : \langle \Gamma; \Delta \rangle$.

Proof. We define a relation R as:

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

Assume $([P_0]_{\alpha'} \mid M_0, \alpha' : \langle \Gamma_0, \Delta_0 \rangle) \in R$:

1. For an input ℓ , because $M_0 = \alpha' : \langle \Gamma_0, \Delta_0 \rangle$ by assumption, that $\alpha' : \langle \Gamma_0, \Delta_0 \rangle \xrightarrow{\ell} \alpha' : \langle \Gamma'_0, \Delta'_0 \rangle$ and $[P_0]_{\alpha'} \mid M_0$ having an input at $\text{subj}(\ell)$ together imply that $[P_0]_{\alpha'} \xrightarrow{\ell} [P'_0]_{\alpha'}$, thus by Lemma 5.1, we have $[P_0]_{\alpha'} \mid \alpha' : \langle \Gamma_0, \Delta_0 \rangle \xrightarrow{\ell} [P'_0]_{\alpha'} \mid M'_0$, and $M'_0 = \alpha' : \langle \Gamma'_0, \Delta'_0 \rangle$. Thus we have $([P'_0]_{\alpha'} \mid M'_0, \alpha' : \langle \Gamma'_0, \Delta'_0 \rangle) \in R$.
2. For an output ℓ , $[P_0]_{\alpha'} \mid M_0 \xrightarrow{\ell} [P'_0]_{\alpha'} \mid M'_0$ implies $M_0 = \alpha' : \langle \Gamma_0, \Delta_0 \rangle \xrightarrow{\ell} \alpha' : \langle \Gamma'_0, \Delta'_0 \rangle = M'_0$. Thus we have $([P'_0]_{\alpha'} \mid M'_0, \alpha' : \langle \Gamma'_0, \Delta'_0 \rangle) \in R$.
3. For an τ , $[P_0]_{\alpha'} \mid M_0 \xrightarrow{\tau} [P_0]_{\alpha'} \mid M_0$ implies that $M_0 = \alpha' : \langle \Gamma_0, \Delta_0 \rangle \xrightarrow{\tau} \alpha' : \langle \Gamma_0, \Delta_0 \rangle = M_0$.

Therefore, by Definition 4.5, R is a satisfaction relation and $\models [P]_\alpha \mid M \triangleright \alpha : \langle \Gamma; \Delta \rangle$ with $M = \alpha : \langle \Gamma; \Delta \rangle$. \square

We define a safety property for partial networks that may include multiple principals. It describes the fact that a monitored network satisfies its specification.

1
2
3
4
5 **Definition 5.3 (Network global safety).** $M \mid M$ is globally safe with respect to
6 Σ if and only if $\models M \mid M \triangleright \Sigma$.
7

8 We introduce a condition on the structure of a network and on its monitors,
9 which guarantees global safety. A partial network is *fully monitored w.r.t.* Σ
10 when all its principals are monitored and the collection of the monitors is weakly
11 bisimilar to Σ . Formally, $M \mid M$ is fully monitored w.r.t. Σ when $M \mid M \equiv$
12 $[P_1]_{\alpha_1} \mid M_1 \mid \dots \mid [P_n]_{\alpha_n} \mid M_n$ for some $n \geq 0$ and $M_1, \dots, M_n \approx \Sigma$. By Theorem
13 5.4 a *fully monitored* network is globally safe. Theorem 5.4 justifies monitoring
14 by ensuring that fully monitored systems behave as expected.
15
16

17 **Theorem 5.4 (Global safety).** If $M \mid M$ is fully monitored w.r.t. Σ , then $\models M \mid$
18 $M \triangleright \Sigma$.
19

20 *Proof.* Assume N is composed by monitored endpoints $[P_i]_{\alpha_i} \mid M_i, i \in \{1, \dots, n\}$.
21

$$M \mid M \equiv [P_1]_{\alpha_1} \mid M_1 \mid \dots \mid [P_n]_{\alpha_n} \mid M_n$$

22 where $M_i = \alpha_i : \langle \Gamma_i; \Delta_i \rangle$ for $i = \{1, \dots, n\}$, $\Sigma = M_1, \dots, M_n$. Based on Theorem
23 5.2, for each $i \in \{1, \dots, n\}$,

$$\models [P_i]_{\alpha_i} \mid M_i \triangleright \alpha_i : \langle \Gamma_i; \Delta_i \rangle$$

24 with $M_i = \alpha_i : \langle \Gamma_i; \Delta_i \rangle$. By Definition 4.5 and induction, we have
25

$$[P_1]_{\alpha_1} \mid M_1 \mid \dots \mid [P_n]_{\alpha_n} \mid M_n \triangleright \alpha_1 : \langle \Gamma_1; \Delta_1 \rangle, \dots, \alpha_n : \langle \Gamma_n; \Delta_n \rangle$$

26 so that $\models M \mid M \triangleright \Sigma$. □
27
28
29
30
31
32
33
34
35
36

37 6. Transparency of Monitored Networks

38 Whereas safety assurance focuses on preventing violations from the principals,
39 transparency ensures that monitors do not affect the behaviour of well-behaved
40 principals. We first consider transparency for partial networks consisting of one
41 single principal (*local transparency*) and then extend the result to monitored net-
42 works with global transport.
43
44
45

46 **Theorem 6.1 (Local transparency).** If $\models [P]_{\alpha} \triangleright \alpha : \langle \Gamma; \Delta \rangle$, then $[P]_{\alpha} \approx ([P]_{\alpha} \mid$
47 $M)$ with $M = \alpha : \langle \Gamma; \Delta \rangle$.
48

49 That is, a correct participant is not impaired by monitoring.
50

51 *Proof.* Define a relation R as:

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

52 Assume $([P]_{\alpha}, [P]_{\alpha} \mid M) \in R$,
53
54
55
56
57

- 1
2
3
4
5
648 • for an output ℓ (the case for τ is similar), $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$ implies $M \xrightarrow{\ell} M'$
6
649 due to $\models [P]_\alpha \triangleright M$; by Lemma 5.1, we have $[P]_\alpha \mid M \xrightarrow{\ell} [P']_\alpha \mid M'$;
7
8
9
650 • for an input ℓ , $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$ only when $M \xrightarrow{\ell} M'$, which together imply
651 that, by Lemma 5.1, $[P]_\alpha \mid M \xrightarrow{\ell} [P']_\alpha \mid M'$.

652 By Definition 4.5, we have $\models [P']_\alpha \triangleright M'$, so that $([P']_\alpha, [P']_\alpha \mid M') \in R$.

653
654 Symmetrically, since, by Theorem 5.2, we have $\models [P]_\alpha \mid M \triangleright \alpha : \langle \Gamma; \Delta \rangle$ with
655 $M = \alpha : \langle \Gamma; \Delta \rangle$,

- 656 • for an output ℓ or τ , $[P]_\alpha \mid M \xrightarrow{\ell} [P']_\alpha \mid M'$ implies $M \xrightarrow{\ell} M'$ whenever
657 $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$;
658 • for an input ℓ , $[P]_\alpha \mid M \xrightarrow{\ell} [P']_\alpha \mid M'$ says $M \xrightarrow{\ell} M'$, which implies
659 $[P]_\alpha \xrightarrow{\ell} [P']_\alpha$.

660 By Definition 4.5, we have $\models [P']_\alpha \mid M' \triangleright M'$, so that $([P']_\alpha \mid M', [P']_\alpha) \in R$. By
661 Definition 4.1, $[P]_\alpha \approx ([P]_\alpha \mid M)$ with $M = \alpha : \langle \Gamma; \Delta \rangle$. \square

662 By Proposition 4.4 and Theorem 6.1, we derive Corollary 6.2 stating that
663 weakly bisimilar static networks combined with the same global transport are
664 weakly bisimilar; i.e. monitoring does not affect routing of information to and
665 from a correct principal.

666 **Corollary 6.2 (Bisimilarity).** If $\models [P]_\alpha \triangleright \alpha : \langle \Gamma; \Delta \rangle$, then for any $\langle r ; h \rangle$, we
667 have $([P]_\alpha \mid \langle r ; h \rangle) \approx ([P]_\alpha \mid M \mid \langle r ; h \rangle)$ with $M = \alpha : \langle \Gamma; \Delta \rangle$.

668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965

$$\langle r ; h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l_j \langle v \rangle} \langle r ; h \cdot s \langle \mathbf{r}_1, \mathbf{r}_2, l_j \langle v \rangle \rangle \rangle$$

$$\begin{array}{l}
1 \\
2 \\
3 \\
4 \\
5 \quad \{\text{REQ}\} \quad \langle r ; h \rangle \xrightarrow{\bar{a}\langle s[\mathbf{r}] : T \rangle}_{\mathbf{g}} \langle r ; h \cdot \bar{a}\langle s[\mathbf{r}] : T \rangle \rangle \\
6 \\
7 \\
8 \quad \{\text{ACC}\} \quad \langle r ; \bar{a}\langle s[\mathbf{r}] : T \rangle \cdot h \rangle \xrightarrow{a\langle s[\mathbf{r}] : T \rangle}_{\mathbf{g}} \langle r ; h \rangle \\
9 \\
10 \\
11 \quad \{\text{SEL}\} \quad \langle r ; h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle}_{\mathbf{g}} \langle r ; h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l\langle v \rangle \rangle \rangle \\
12 \\
13 \\
14 \quad \{\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}_{\mathbf{g}} \langle r ; h \rangle \\
15 \\
16 \\
17 \quad \{\text{NET}\} \quad \frac{N = [P]_{\alpha} | \langle r ; h \rangle \quad [P]_{\alpha} \xrightarrow{\ell} [P']_{\alpha} \quad \langle r ; h \rangle \xrightarrow{\ell}_{\mathbf{g}} \langle r' ; h' \rangle \quad N' = [P']_{\alpha} | \langle r' ; h' \rangle}{N \xrightarrow{\ell}_{\mathbf{g}} N'} \\
18 \\
19 \\
20 \\
21 \quad \{\text{TAU}\} \quad \frac{N \longrightarrow N'}{N \xrightarrow{\tau}_{\mathbf{g}} N'} \quad \{\text{RES}\} \quad \frac{N \xrightarrow{\ell}_{\mathbf{g}} N' \quad a \notin \text{subj}(\ell)}{(\nu a)N \xrightarrow{\ell \setminus a}_{\mathbf{g}} (\nu a)N'} \quad \{\text{STR}\} \quad \frac{N \equiv N_0 \xrightarrow{\ell}_{\mathbf{g}} N'_0 \equiv M'}{N \xrightarrow{\ell}_{\mathbf{g}} N'} \\
22 \\
23 \\
24 \quad \{\text{PAR}\} \quad \frac{N_1 \xrightarrow{\ell}_{\mathbf{g}} N'_1 \quad \text{bn}(\ell) \cap \text{fn}(N_2) = \emptyset \quad \text{dest}(\ell) \notin \mathcal{P}(N_2)}{N_1 \parallel N_2 \xrightarrow{\ell}_{\mathbf{g}} N'_1 \mid N_2} \\
25 \\
26 \\
27 \\
28 \quad \{\text{MON}\} \quad \frac{\mathbf{N} = N \mid M \quad N \xrightarrow{\ell}_{\mathbf{g}} N' \quad M \xrightarrow{\ell} M' \quad N' = N' \mid M'}{\mathbf{N} \xrightarrow{\ell}_{\mathbf{g}} N'} \\
29 \\
30 \\
31 \\
32 \\
33 \\
34 \\
35 \\
36 \\
37 \\
38 \\
39 \\
40 \\
41 \\
42 \\
43 \\
44 \\
45 \\
46 \\
47 \\
48 \\
49 \\
50 \\
51 \\
52 \\
53 \\
54 \\
55 \\
56 \\
57 \\
58 \\
59 \\
60 \\
61 \\
62 \\
63 \\
64 \\
65
\end{array}$$

Figure 10: LTS for networks

Similarly, the parallel composition of a principal sending $s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle$ and the global transport is made visible as follows:

$$[s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle; P']_{\alpha} | \langle r ; h \rangle \xrightarrow{s[\mathbf{r}_1, \mathbf{r}_2]!l_j\langle v \rangle}_{\mathbf{g}} [P']_{\alpha} | h \cdot s\langle \mathbf{r}_1, \mathbf{r}_2, l_j\langle v \rangle \rangle$$

668 We define dest as a partial function mapping a label, which representing an
669 action, to its destination as:

$$\begin{aligned}
\text{dest} ::= & \bar{a}\langle s[\mathbf{r}] : T \rangle \mapsto a \mid a\langle s[\mathbf{r}] : T \rangle \mapsto a \\
& \mid s[\mathbf{r}_1, \mathbf{r}_2]!l\langle v \rangle \mapsto s[\mathbf{r}_2] \mid s[\mathbf{r}_1, \mathbf{r}_2]?l\langle v \rangle \mapsto s[\mathbf{r}_2]
\end{aligned}$$

670 The notation of global observable transition $\xrightarrow{\ell}_{\mathbf{g}}$, used to denote *globally observable*
671 *action* ℓ , is defined by the rules in Fig. 10. Rules $\{\text{REQ}\}$ and $\{\text{ACC}\}$ (resp.
672 $\{\text{SEL}\}$ and $\{\text{BRA}\}$) are for inserting and removing invitation messages (resp. mes-
673 sages in established sessions) from the global transport. Rules $\{\text{ACC}\}$ and $\{\text{BRA}\}$
674 represent that, as a message leaves the global queue, there should be a local principal
675 receiving it as an input. Similarly, rules $\{\text{REQ}\}$ and $\{\text{SEL}\}$ represent that, as

1
2
3
4
5 a message enters the global queue, there should be a local principal outputting
6 it to the queue. By $\{\text{NET}\}$ for unmonitored networks, as $N \xrightarrow{\ell}_g N'$, it means
7 $\exists [P]_\alpha \in N, [P]_\alpha \xrightarrow{\ell} [P']_\alpha$ (i.e. locally visible) such that $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$ (i.e.
8 globally visible). Rule $\{\text{TAU}\}$ summarizes the reduction rules defined in Section
9 3. Rule $\{\text{RES}\}$ and $\{\text{STR}\}$ are standard. Rule $\{\text{PAR}\}$ says that, the bound names of
10 action ℓ should not be any free name appearing in network N_2 , and it should not be
11 absorbed by any process in network N_2 (i.e. its destination is not in N_2). By rule
12 $\{\text{MON}\}$ for monitored networks, $N \xrightarrow{\ell}_g N'$ means $\exists [P]_\alpha \mid M \in N, [P]_\alpha \xrightarrow{\ell} [P']_\alpha$
13 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
14 visible).
15
16
17
18

19 **Theorem 6.3 (Global transparency).** Assume N and N' have the same global
20 transport $\langle r ; h \rangle$. If N is fully monitored w.r.t. Σ and $N' = M \mid \langle r ; h \rangle$ is
21 unmonitored but $\models M \triangleright \Sigma$, then we have $N \approx N'$.
22
23

24 *Proof.* Define a relation R :

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

25
26
27
28
29 We prove that R is a standard strong bisimilar relation over $\xrightarrow{\ell}_g$. Note that, $M \triangleright \Sigma$
30 means $\forall [P_i]_{\alpha_i} \in M$, we have $\alpha_i : \langle \Gamma_i ; \Delta_i \rangle \in \Sigma$ and $\models [P_i]_{\alpha_i} : \alpha_i : \langle \Gamma_i ; \Delta_i \rangle$.
31

32 1. As $N \xrightarrow{\ell}_g N'$, it implies $\exists [P_j]_{\alpha_j} \mid M_j \in N, [P_j]_{\alpha_j} \xrightarrow{\ell} [P'_j]_{\alpha_j}$ and $M_j \xrightarrow{\ell} M'_j$
33 such that $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$, and other monitored processes in N are
34 not affected. When ℓ is an input, by Definition 4.5, since $\models M \triangleright \Sigma$, we
35 should have $[P_j]_{\alpha_j} \xrightarrow{\ell} [P'_j]_{\alpha_j}$; when ℓ is an output or a τ action, by Def-
36 inition 4.5, the transition of $[P_j]_{\alpha_j} \xrightarrow{\ell} [P'_j]_{\alpha_j}$ is able to take place. Both
37 cases lead to $M \xrightarrow{\ell} M'$ and $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$ so that $N' = M' \mid$
38 $\langle r ; h \rangle \xrightarrow{\ell}_g M' \mid \langle r' ; h' \rangle = N'$, and $\models [P'_j]_{\alpha_j} : \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle$ by
39 Definition 4.5. $\alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle$ is the resulting new configuration of α_j in
40 Σ . Other specifications $\{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I \setminus \{j\}} \in \Sigma$ are not affected. Let
41 $\Sigma' = \alpha_j : \langle \Gamma'_j ; \Delta'_j \rangle, \{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I \setminus \{j\}}$. Therefore, for the resulting
42 new network $N' = M' \mid \langle r' ; h' \rangle$, we have $\models M' \triangleright \Sigma'$. Thus we have
43 $(N, N') \in R$.
44
45
46
47
48
49

50 2. For the symmetric case, as $N \xrightarrow{\ell}_g N'$, it implies $\exists [P_j]_{\alpha_j} \in N, [P_j]_{\alpha_j} \xrightarrow{\ell}$
51 $[P'_j]_{\alpha_j}$ such that $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$ and other processes in N are not af-
52 fected. Since $\models M \triangleright \Sigma$, without loss of generality, let $M_j = \alpha_j : \langle \Gamma_j ; \Delta_j \rangle$,
53 then we have, for any ℓ , $[P_j]_{\alpha_j} \mid M_j \xrightarrow{\ell} [P'_j]_{\alpha_j} \mid M'_j$, where $M'_j = \alpha_j :$
54 $\langle \Gamma'_j ; \Delta'_j \rangle$. It makes $\langle r ; h \rangle \xrightarrow{\ell}_g \langle r' ; h' \rangle$, so that $N \xrightarrow{\ell}_g N'$. Since N'
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5 708 is a fully monitored network, its static part (i.e. the part when the global
6 709 transport is taken off from N'), say $[P_i]_{\alpha_i} \mid \{M_i\}_{i \in I}$ where $\{M_i\}_{i \in I} =$
7 $\alpha_j : \langle \Gamma'_j; \Delta'_j \rangle, \{\alpha_i : \langle \Gamma_i; \Delta_i \rangle\}_{i \in I \setminus \{j\}}, \models [P_i]_{\alpha_i} \mid \{M_i\}_{i \in I} : \Sigma'$ where
8 710 $\Sigma' = \alpha_j : \langle \Gamma'_j; \Delta'_j \rangle, \{\alpha_i : \langle \Gamma_i; \Delta_i \rangle\}_{i \in I \setminus \{j\}}$. Thus we have $(N', N') \in R$.
9 711 \square
10 712

11
12 713 By Theorems 5.4 and 6.3, we can *mix* unmonitored principals with monitored
13 714 principals still obtaining global safety assurance:

15 715 **Corollary 6.4 (Mixed Network).** If $M \mid M$ is fully monitored with respect to Σ ,
16 716 $\models M' \triangleright \Sigma$, and $\mathcal{P}(M) \cap \mathcal{P}(M') = \emptyset$, then $\models (M \mid M) \mid M' \triangleright \Sigma$.

17
18
19 717 In the above corollary, untyped M is monitored by M which specifies Σ , while
20 718 M' is unmonitored but statically checked to conform to Σ . The result shows that
21 719 they can safely be composed.

22 23 24 720 7. Session Fidelity

25
26 721 The property of session fidelity says that, whenever all the principals in a
27 722 static network conform to their specifications, then all of the derivatives of this
28 723 static network conform to evolutions of the initial global specification.

29
30
31 724 **Global Safety vs Session Fidelity.** Recall that global safety (Definition 5.3) only
32 725 ensures that in a network where principals are well-behaved with respect to their
33 726 local types, all interactions conform to the collection of these local types. Session
34 727 fidelity is a stronger property than global safety (Definition 5.3) as illustrated in
35 728 Example 7.1.

36
37
38 **Example 7.1.** Consider a simple global type

$$39 G = \mathbf{r}_1 \rightarrow \mathbf{r}_2 : \{l_1(x_1)\{x_1 > 9\}.G_1, l_2(x_2)\{x_2 < 10\}.G_2\}$$

40
41 and processes P and Q implementing roles \mathbf{r}_1 and \mathbf{r}_2 in established session s

$$42 P = s[\mathbf{r}_1, \mathbf{r}_2]!l_1\langle 10 \rangle.P'$$

$$43 Q = s[\mathbf{r}_1, \mathbf{r}_2]? \{l_i(x_i).Q_i\}_{i \in \{1,2\}}$$

44
45 729 Suppose that during runtime P sends out message $s\langle \mathbf{r}_1, \mathbf{r}_2, l_1\langle 10 \rangle \rangle$ but, Q receives
46 730 a message, perhaps revised by an attack, $s\langle \mathbf{r}_1, \mathbf{r}_2, l_2\langle 8 \rangle \rangle$. These actions satisfy
47 731 global safety since satisfy the specifications of P and Q , namely they are *locally*
48 732 well-behaved. This scenario (i.e., the content of the message being modified be-
49 733 tween a send and a corresponding receive action) does not conform to the intended
50 734 global protocol. We define a property, session fidelity, that rules out the scenario
51 735 above (Definition 7.8) and prove (Theorem 7.13) that fully monitored networks

1
2
3
4
5 736 with global transport satisfy session fidelity. This is due to the fact that: (1) s is a
6 737 private session ID that can be viewed only by the participants in the session, (2) all
7 738 principals are guarded by monitors hence all messages reaching the global trans-
8 739 port are valid, and (3) the global transport preserves the values of the messages it
9 740 gets.

11
12 741 **Configurations.** We define session fidelity after giving the labelled transition re-
13 742 lation for *configurations*, and a few auxiliary definitions.

14
15 743 **Definition 7.2 (Configuration).** A configuration is denoted by $\Phi = \Sigma; \langle r ; h \rangle$,
16 744 where all messages corresponding to the actions guarded by Σ are in h .

17
18 745 A configuration guides the global behaviours in a network. By including the
19 746 global queue, we let configuration capture the *global* behaviour in a network,
20 747 which accounts also for the correct routing and dispatch of messages. Before
21 748 giving the semantics of configurations, it will be useful to define when and how
22 749 configurations can be composed. Let $\mathcal{P}(\Phi)$ be the set of principals involving in Φ .
23 750

24
25
26 751 **Definition 7.3 (Parallel composition of configurations).** Let $\Phi_1 = \Sigma_1; \langle r_1 ; h_1 \rangle$
27 752 and $\Phi_2 = \Sigma_2; \langle r_2 ; h_2 \rangle$ be configurations. We say that Φ_1 and Φ_2 are *composable*
28 753 whenever $\mathcal{P}(\Phi_1) \cap \mathcal{P}(\Phi_2) = \emptyset$ and the union of their routing tables remains a
29 754 function. If Φ_1 and Φ_2 are composable, then we define the composition of Φ_1 and
30 755 Φ_2 as: $\Phi_1 \odot \Phi_2 = \Sigma_1, \Sigma_2; \langle r_1 \cup r_2 ; h_1 \cdot h_2 \rangle$.

31
32 756 The formal semantics of configurations is defined by the LTS in Fig. 11. The
33 757 behaviour of each principal in a network is guided by the specification Σ , and is
34 758 observed by the global transport $\langle r ; h \rangle$. Except rules [Acc] and [Par], all rules
35 759 are straightforward from the LTS of specifications (defined in Section 4.1) and the
36 760 one of dynamic networks (Fig. 10). We comment below on the interesting rules.

- 37
38
39 761 1. Rule [Acc] indicates that, only when the invitation has been (internally)
40 762 accepted by a principal in the network, the routing information registers
41 763 $s[r] \mapsto \alpha$. When we observe the global transport (externally), we only
42 764 observe that an invitation is moved out from the global queue (which implies
43 765 that it has been accepted). However, we do not know *who* accepts it. Only
44 766 Σ tells which principal accepts this invitation, so that we can register it in
45 767 the routing information using α .
46
47 768 2. Rule [Par] says if Φ_1 and Φ_2 are composable (Definition 7.3), after Φ_1 be-
48 769 comes as Φ_2 , they are still composable.

49
50
51 770 Our framework relies on two structural (well-formedness) properties on speci-
52 771 fications: *consistency* and *coherence*. Consistent specifications are the ones corre-
53 772 sponding to well-formed concrete systems (i.e., where the session initiation proce-
54 773 dures are well-regulated, and where the active sessions correspond to projections
55 774 of some well-formed global type).

$$\begin{array}{l}
1 \\
2 \\
3 \\
4 \\
5 \\
6 \\
7 \\
8 \\
9 \\
10 \\
11 \\
12 \\
13 \\
14 \\
15 \\
16 \\
17 \\
18 \\
19 \\
20 \\
21 \\
22 \\
23 \\
24 \\
25 \\
26 \\
27 \\
28 \\
29 \\
30 \\
31 \\
32 \\
33 \\
34 \\
35 \\
36 \\
37 \\
38 \\
39 \\
40 \\
41 \\
42 \\
43 \\
44 \\
45 \\
46 \\
47 \\
48 \\
49 \\
50 \\
51 \\
52 \\
53 \\
54 \\
55 \\
56 \\
57 \\
58 \\
59 \\
60 \\
61 \\
62 \\
63 \\
64 \\
65
\end{array}$$

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

Figure 11: Labelled transition relation for configurations

Definition 7.4 (Consistent and coherent specifications). $\Sigma = \{\alpha_i : \langle \Gamma_i ; \Delta_i \rangle\}_{i \in I}$ 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 $a : 0(T[\mathbf{r}])$ 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.

Next, we define *receivability*, *configurational consistency* and *conformance* for configurations, which are based on the LTS of configurations and dynamic networks. Receivability entices the ability for a message in transit to reach its destination.

Definition 7.5 (Receivable configuration). Receivability of a configuration $\Sigma ; \langle r ; h \rangle$ is defined 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 $\Sigma ; \langle r ; m \cdot h' \rangle \xrightarrow{\ell}_{\mathbf{g}} \Sigma' ; \langle r' ; h' \rangle$, where ℓ corresponding to m , and $\Sigma' ; \langle r' ; h' \rangle$ is receivable.

1
2
3
4
5 793 A configuration $\Sigma; \langle r ; h \rangle$ is *configurationally consistent* if all of its multi-step
6 794 global input transition derivatives can be performed and the resulting specifica-
7
8 795 tions Σ is consistent (according to Definition 7.4).

9
10 796 **Definition 7.6 (Configurational consistency).** A configuration $\Phi = \Sigma; \langle r ; h \rangle$ is
11 797 *configurationally consistent* whenever

- 12
13 798 1. h is empty and Σ is consistent, or
14 799 2. h is not empty, $\Sigma; \langle r ; h \rangle$ is receivable, and after receiving all messages in
15
16 800 h with $\Sigma \xrightarrow{\ell_1 \dots \ell_n} \Sigma'$ (by the LTS in Figure 7), where $\ell_i, i = \{1, \dots, n\}$ are
17 801 inputs and, $\forall m \in h, \exists \ell \in \ell_1 \dots \ell_n$ such that ℓ corresponds to m , we have
18 802 Σ' is consistent.

19
20
21 803 In other words, $\Sigma; \langle r ; h \rangle$ is configurationally consistent if, in each of its deriva-
22 804 tives, all messages in the transport can be “received” by some monitors in Σ and,
23 805 after absorbing all these messages, the resulting Σ' is still consistent. Confor-
24 806 mance links networks and configurations.

25
26
27 807 **Definition 7.7 (Conformance to a configuration).** Assume a network $N \equiv M \mid$
28 808 $\langle r ; h \rangle$ is given. We say that N conforms to $\Sigma; \langle r ; h \rangle$ when:

- 29
30 809 1. h is empty, $\models M \triangleright \Sigma$ and Σ is consistent, or
31 810 2. h is not empty, and the following conditions hold
32
33 811 (a) $\models M \triangleright \Sigma$,
34 812 (b) all messages in h are receivable to M , and
35 813 (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
36 814 each $\ell_i, i = \{1, \dots, n\}$ is an input, Σ' is consistent.

37
38
39 815 **Session Fidelity.** Session fidelity describes the relation between a network and
40 816 the configuration specifying it: all evolutions of the network should correspond to
41 817 expected evolutions of the configuration which does not lead to ill-formed config-
42 818 urations. We now give the formal definition of session fidelity.

43
44
45 819 **Definition 7.8 (Session fidelity).** Assume configuration $\Sigma; \langle r ; h \rangle$ is configura-
46 820 tionally consistent. We say that N satisfies session fidelity w.r.t. $\Sigma; \langle r ; h \rangle$ if and
47 821 only if, for any $\ell, N \xrightarrow{\ell}_g N'$ implies $\Sigma; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r' ; h' \rangle$ and $\Sigma'; \langle r' ; h' \rangle$ is
48 822 configurationally consistent and N' satisfies session fidelity w.r.t. $\Sigma'; \langle r' ; h' \rangle$.

49
50
51 823 Before proving session fidelity for our monitored framework we give a few
52 824 auxiliary lemmas. Lemma 7.9 states that, as a network conforms to some configu-
53 825 rationally consistent configuration, the evolution of the configuration must be able
54 826 to consume an output occurrence in the network:

1
2
3
4
5 827 **Lemma 7.9.** Assume a network $N \equiv M | \langle r ; h \rangle$ conforms to $\Sigma ; \langle r ; h \rangle$, and that
6
7 828 $\Sigma ; \langle r ; h \rangle$ is configurationally consistent. If $N \xrightarrow{\ell}_g N'$ with ℓ being an output and
8
9 829 $\Sigma ; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma' ; \langle r ; h \cdot m \rangle$, then $\Sigma' ; \langle r ; h \cdot m \rangle$ is receivable.

10
11 830 *Proof.* We only show the interesting case. When $\ell = \bar{a} \langle s[r] : T \rangle$, since Σ is
12
13 831 consistent, by Definitions 7.4, there exists $a : \mathbb{I}(T[r])$ in some Γ of Σ . Because ℓ
14
15 832 does not affect the existence of $a : \mathbb{I}(T[r])$, it remains in Γ of Σ' , thus invitation
16
17 833 $m = \bar{a} \langle s[r] : T \rangle$ is receivable to Σ' .

Let $\alpha_i = \langle \Gamma_i, \Delta_i \rangle$. When $\ell = s[r_1, r_2] ! l_j \langle v \rangle$, by Definitions 7.4 and 7.7, since
18
19 $\models M \triangleright \Sigma$ and Σ is consistent, $\exists \alpha_s, \alpha_r \in \Sigma$, $\exists G$ is well-formed of the form

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

20
21 834 such that s obeys to G :

$$\begin{aligned} \Delta_s(s[r_1]) &= G \upharpoonright \mathbf{r}_1 = \mathbf{r}_2 ! \{l_i(x_i : (T[r])_i) \{A_i\}. G_i \upharpoonright \mathbf{r}_1\}_{i \in I} \\ \Delta_r(s[r_2]) &= G \upharpoonright \mathbf{r}_2 = \mathbf{r}_1 ? \{l_i(x_i : (T[r])_i) \{A'_i\}. G_i \upharpoonright \mathbf{r}_2\}_{i \in I} \end{aligned} \quad (1)$$

22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

As action $s[r_1, r_2] ! l_j \langle v \rangle$ fires, Equation 1 changes to

$$\begin{aligned} \Delta_s(s[r_1]) &= G_j \upharpoonright \mathbf{r}_1 \\ \Delta_r(s[r_2]) &= G \upharpoonright \mathbf{r}_2 = \mathbf{r}_1 ? \{l_i(x_i : (T[r])_i) \{A'_i\}. G_i \upharpoonright \mathbf{r}_2\}_{i \in I} \end{aligned}$$

835 the receiving capability of $\mathbf{r}_1 ?$ still remains in $\Delta_r(s[r_2])$, where $\alpha_r \in \Sigma'$, thus
836 $m = s \langle \mathbf{r}_1, \mathbf{r}_2, l_j \langle v \rangle \rangle$ is receivable to Σ' .

837 As $N \equiv M | H$ and $\models M \triangleright \Sigma$, the satisfaction relation of M and Σ remains
838 whenever action takes place.

839 Lemma 7.10 says that, if the static part of a network satisfies a specification,
840 then the evolution of the static part still satisfies the corresponding evolution of
841 the specification.

842 **Lemma 7.10.** Assume $N \equiv M | H$ and $\models M \triangleright \Sigma$. If $N \xrightarrow{\ell}_g N' \equiv M' | H'$ and
843 $\Sigma \xrightarrow{\ell} \Sigma'$, then $\models M' \triangleright \Sigma'$.

844 *Proof.* Directly from Definition 4.5.

845 Finally, Lemma 7.12 states that, if a network conforms to a configurationally
846 consistent configuration, then any evolution of the network conforms to the corre-
847 sponding evolution of the configuration, which is still configurationally consistent.
848 Lemma 7.12 relies on the definition of routing table given below.

Definition 7.11 (Routing table). . We define $\text{route}(\Sigma)$, the routing table derived from Σ , as follows:

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

849 The routing table is used to observe inputs. Note that by Definition 7.4 (2), as long
850 as Σ is consistent, the existence of $a : \mathbf{O}(T[\mathbf{r}])$ in Γ implies that the corresponding
851 $a : \mathbf{I}(T[\mathbf{r}])$ is also in Γ .

852 **Lemma 7.12.** Assume configuration $\Sigma; \langle r ; h \rangle$ is configurationally consistent,
853 and network $N \equiv M \langle r ; h \rangle$ conforms to configuration $\Sigma; \langle r ; h \rangle$. Then for any
854 ℓ , whenever we have $N \xrightarrow{\ell}_g N'$ such that $\Sigma; \langle r ; h \rangle \xrightarrow{\ell}_g \Sigma'; \langle r' ; h' \rangle$, it holds that
855 $\Sigma'; \langle r' ; h' \rangle$ is configurationally consistent and that N' conforms to $\Sigma'; \langle r' ; h' \rangle$.

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

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

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

863 Since Σ allows ℓ and Σ is consistent, then $\exists \alpha_r, \alpha_s \in \Sigma$, and $\exists G$ well-formed
864 of the form

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

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

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

864 Case 1: h is empty. By Lemma 7.9, after receiving m , say $\Sigma' \xrightarrow{\ell} \Sigma''$, Σ''
865 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 Defini-
866 tion 7.4. By Definition 7.6, $\Sigma'; \langle r ; m \rangle$ is configurationally consistent, and

867 $\models M' \triangleright \Sigma'$ by Lemma 7.10, thus N' conforms to $\Sigma'; \langle r ; h \cdot m \rangle$.

868

869 **Case 2:** h is not empty. Since $\Sigma; \langle r ; h \rangle$ is configurationally consistent,
 870 again, by Lemma 7.9, after receiving messages in h (but not m), say $\Sigma' \xrightarrow{\ell_0 \dots \ell_n}$
 871 Σ'_1 , where every action in $\ell_0 \dots \ell_n$ corresponds to each message in h , we
 872 have $\Sigma'_1; \langle r' ; m \rangle$ is configurationally consistent. After Σ'_1 receives m , say
 873 $\Sigma'_1 \xrightarrow{s[p_1, p_2]?l\langle v \rangle} \Sigma''$, where $s[p_1, p_2]?l\langle v \rangle$ is dual to ℓ , with the same reasoning
 874 above, Σ'' has $s[r_1] = G'_j \upharpoonright r_1$ and $s[r_2] = G'_j \upharpoonright r_2$, so that Σ'' is consis-
 875 tent. By Definition 7.6, $\Sigma'; \langle r ; h \cdot m \rangle$ is configurationally consistent, and
 876 $\models M' \triangleright \Sigma'$ by Lemma 7.10, thus N' conforms to $\Sigma'; \langle r ; h \cdot m \rangle$.

877

878 **(Bra)** Let $\ell = s[r_1, r_2]?l_j\langle v \rangle$, $N \xrightarrow{\ell}_g N'$ and N conforms to $\Sigma; \langle \text{route}(\Sigma) ; h \rangle$.

879

880 **Case 1:** h is empty. Since $\Sigma; \langle \text{route}(\Sigma) ; \emptyset \rangle \not\xrightarrow{\ell}_g$, so this case never happens.

881

882 **Case 2:** h is not empty. Thus, $N \xrightarrow{\ell}_g N'$ and

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

882 where h/m means taking off message m from h , where $m = s\langle r_1, r_2, l_j\langle v \rangle \rangle$

883

884 We have $r = \text{route}(\Sigma) = \text{route}(\Sigma')$ because there is no change to the el-
 885 elements in Σ or to the routing table. By Definition 7.6, after receiving all
 886 messages in H , Σ is consistent, thus Σ' , which has received message m is
 887 consistent after receiving all messages in h/m . By Lemma 7.10, we have
 888 $\models M' \triangleright \Sigma'$ thus N' conforms to $\Sigma'; \langle r ; h/m \rangle$.

889

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

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

890 where $m = \bar{a}\langle s[r] : T \rangle$. Then $r = \text{route}(\Sigma) = \text{route}(\Sigma')$ because, by Defi-
 891 nition 7.11, nothing new is registered to the routing table.

892

893 Since Σ allows ℓ and Σ is consistent, by Definition 7.4, $\exists \Gamma_i, \Gamma_j \in \Sigma$ such
 894 that $a : \mathbf{I}(T[r]) \in \Gamma_i$ and $a : \mathbf{O}(T[r]) \in \Gamma_j$. After $\Sigma \xrightarrow{\ell} \Sigma'$, by rule [REQ] in
 895 the LTS of specifications, $a : \mathbf{I}(T[r])$ remains in Γ'_i , $a : \mathbf{O}(T[r])$ remains in

1
2
3
4
5 896 Γ'_j , and thus they both remain in Σ' .

6 897
7
8 898 Case 1: h is empty. By Lemma 7.9, after receiving m , say $\Sigma' \xrightarrow{a\langle s[r]:T \rangle} \Sigma''$,
9 899 both $a : \mathbb{I}(T[r])$ and $a : \mathbb{O}(T[r])$ remain in Σ'' , satisfying Definition 7.4, so
10 900 that $\Sigma'; \langle r ; m \rangle$ is configurationally consistent. By Lemma 7.10, we have
11 901 $\models M' \triangleright \Sigma'$, thus N' conforms to $\Sigma'; \langle r ; h \cdot m \rangle$.
12
13 902

14
15 903 Case 2: h is not empty. The proof is similar to the one in (Sel) and omitted.
16 904

17
18 905 **(Acc)** Let $\ell = a\langle s[r] : T \rangle$.
19 906

20
21 907 Case 1: h is empty. Since $\Sigma; \langle \text{route}(\Sigma) ; \emptyset \rangle \not\xrightarrow{\ell}_g$, this case never happens.
22 908

23
24 Case 2: h is not empty. If $N \xrightarrow{\ell}_g N'$ and

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

25
26
27
28
29 909 where $m = \bar{a}\langle s[r] : T \rangle$. Since there exists $\Delta \in \Sigma'$ s.t. $s[r] \in \Delta$, by Defini-
30 910 tion 7.11, $r' = \text{route}(\Sigma)$, $s[r] \mapsto \alpha = \text{route}(\Sigma')$.
31 911

32
33 912 For the same reasoning in (Bra), we have $\Sigma'; \langle r ; h/m \rangle$ is configurationally
34 913 consistent. By Lemma 7.10, we have $\models M' \triangleright \Sigma'$ thus N' conforms to
35 914 $\Sigma'; \langle r ; h/m \rangle$.
36 915

37
38
39
40 916 The proof for other cases is trivial.

41
42 917 **Theorem 7.13 (Session fidelity).** If N is fully monitored and conforms to $\Sigma; \langle r ; h \rangle$,
43 918 which is configurationally consistent, then N satisfies session fidelity.

44
45 919 *Proof.* The proof is straightforward by Lemma 7.12 and Definition 7.8.
46

47
48 920 **Proposition 7.14.** Whenever a network is fully monitored, global safety implies
49 921 session fidelity.

50
51 922 *Proof.* Simply by Definitions 7.6 and 7.7 and Corollary 6.2 and Theorems 5.4 and
52 923 7.12.
53
54
55
56
57
58
59
60
61
62
63
64
65

8. Related Work

Monitors. Our work features a located, distributed process calculus to model *dynamic* monitored networks. An account of the state of the art of runtime monitors can be found in [29, 38]. According to Havelund and Goldberg [29], specification-based runtime verification consists of monitoring a program's execution against a user-provided specification of the intended program's behaviour. Leucker and Schallhart [38] define runtime verification as the discipline of dealing with the detection of violations (or satisfaction) of correctness properties. They point out the use of runtime verification for contract enforcement.

Global specification languages. Message Sequence Charts (MSC), which are also known as UML sequence diagrams, have been the focus of many works [29, 36, 27, 37]. Among them, Kruger et al. [37] propose a runtime monitoring framework based on projecting MSC to distributed monitors based on finite state machines. They use aspect-oriented programming techniques to inject the monitors into the implementation of the components. Gan [26] follows the same path, but with a centralised approach. Both works do not provide a formal model, formal guarantees of correctness, nor support behavioural analysis. BPEL [4, 5, 27] is an orchestration description language that is now a common part of many industrial distributed systems where web services must be used in a coordinated manner. It supports the definition of abstract specifications as well as their execution. BPEL specifications are designed to be run in a centralised way. Baresi et. al [4] developed a run-time monitoring tool with assertions based on BPEL as an execution language. When the execution of a BPEL process reaches the point where an assertion must be checked, the tool calls an external service to check its satisfaction. This work does not consider properties, such as transparency and local/global safety. On another line of research, van der Aalst et al. [50] use abstract BPEL process as specifications. Their work focuses on checking conformance between *execution logs* (obtained by observing a number of executions based on SOAP message exchanges, and then translated into Petri Nets) and choreographies expressed as abstract BPEL processes. The focus of [50] is on checking conformance a posteriori, as well as on revealing (mining) and re-engineering choreographies according to the actual system's behaviour. Differently from [50] our work establishes a theory of dynamic monitoring. The aim of our work is to observe communication as they occur to prevent unsafe interactions, while providing a formal framework that complements static (behavioural) typing techniques, and supports reasoning about equivalence of networks. Finally, WS-CDL is a more recent description language which aims at describing decentralised choreographies. Cambronero et al. [11] transform choreographies written in WS-CDL into timed-automata and verify systems against them. The work in [11] does not

1
2
3
4
5 963 develop a projection algorithm nor ensures global conformance with respect to a
6 964 choreography.

7
8
9 965 **Theory of monitored networks.** The work of Ferrari et al. [25] proposes an
10 966 ambient-based run-time monitoring formalism, called guardians, targeted at ac-
11 967 cess control rights for network processes, and Klaim [19] advocates a hybrid
12 968 (dynamic and static) approach for access control against capabilities (policies)
13 969 to support static checking integrated within a dynamic access-control procedure.
14 970 These works address specific forms of access control for mobility, while our more
15 971 general approach aims at ensuring correct behaviour in sessions through a combi-
16 972 nation of static and run-time verification.

17
18
19 973 The work of Capecchi et al. [12] presents a monitor-based information-flow
20 974 analysis of multiparty sessions. The monitors in [12] are inline (following [14])
21 975 and control the information-flow by tagging each message with security levels.
22 976 Since each inlined monitor is located within a local process, the interactions be-
23 977 tween endpoint processes and their corresponding monitors are *synchronous*. We
24 978 study *asynchronous* communications that, while being closer to an actual network
25 979 implementation, introduce considerable challenges in the development of a theory.
26 980 Other works on inlined monitors, such as [28, 1, 49], provide a policy specifica-
27 981 tion language. The aim is to write policies into the monitors, with the guarantee
28 982 that the specifications in the inlined monitors satisfy the original policies. Inline
29 983 monitors require direct access to the code, whereas our approach, outline monitor-
30 984 ing (i.e., the implantation of monitors is independent from the implementation of
31 985 the observed applications), ensures interoperability with any language and archi-
32 986 tecture. Other related works on monitoring conversations are [48, 2]. Simmonds
33 987 et al. [48] propose a runtime monitoring approach based on MSC as a specifica-
34 988 tion language to represent global protocols, and transform MSC specifications into
35 989 automata. They provide conformance checking of *finite* execution traces against
36 990 specifications. Ancona et al. [2, 40] propose a dynamic monitoring framework
37 991 based on MPST for Multi-Agent Systems (MAS) to guard interactions between
38 992 local agents and their environments. They gave a procedure that automatically
39 993 derives a self-monitoring MAS from Jason (a MAS development platform), and
40 994 verifies that a MAS implementation is compliant with a given *global session type*,
41 995 which can naturally be represented as cyclic Prolog terms. Their monitoring is
42 996 only synchronous. Their development focuses on implementation and does not
43 997 involve proofs of formal properties.

44
45
46
47
48
49
50
51 998 **Monitoring and MPST.** An informal approach to monitoring based on MPST,
52 999 and an outline of monitors are presented in [17]. However, [17] only gives an
53 1000 overview of the desired properties, and requires all local processes to be dynami-
54 1001 cally verified through the protections of system monitors. In this article, instead,

1
2
3
4
5 1002 we integrate statically and dynamically verified local processes into one network,
6 1003 and formally state the properties of this combination. Some recent works [3, 18]
7 1004 use multiparty session types for dynamic updates. Anderson et al. [3] study
8 1005 channel conditions of running processes to be able to update them and ensure
9 1006 deadlock-freedom, while a system in Coppo et al. [18] enables to update global
10 1007 types dynamically. The work [18] is based on the formulation (without assertions)
11 1008 studied in this article. Recently, Jia and al. [35] proposed a linear-logic based
12 1009 session-calculus close to ours describing monitor semantics for higher-order ses-
13 1010 sions which include rules for *blame assignment*.

14 1011 In summary, compared to these related works, our contribution focuses on the
15 1012 enforcement of global safety, with protocols specified as multiparty session types
16 1013 with assertions. It also provides formalisms and theorems for decentralised run-
17 1014 time monitoring, targeting interaction between components written in multiple
18 1015 (e.g., statically and dynamically typed) programming languages.

19 1016 **9. Conclusion and future work**

20 1017 We proposed a new formal safety assurance framework to specify and enforce
21 1018 global safety of distributed systems through dynamic verification. We formally
22 1019 proved the correctness of our architectural framework through a π -calculus based
23 1020 theory, identified in two key properties of dynamic networks: global transparency
24 1021 and safety. We introduced a behavioural theory over monitored networks which
25 1022 allows compositional reasoning over trusted and untrusted (but monitored) com-
26 1023 ponents.

27 1024 **Implementations.** As a part of our collaboration with the Ocean Observatories
28 1025 Initiative [44], our theoretical framework is currently realised by an implementa-
29 1026 tion [34, 43, 21], in which each monitor supports all well-formed protocols and
30 1027 is automatically self-configured, via session initiation messages, for all sessions
31 1028 that the endpoint participates in. Our implementation of the framework automates
32 1029 distributed monitoring by generating FSM from the local protocol projections. In
33 1030 this implementation, the global protocol serves as the key abstraction that helps
34 1031 unify the aspects of specification, implementation and verification (both static and
35 1032 dynamic) of distributed application development. Our experience has shown that
36 1033 the specification framework can accommodate diverse practical use cases, includ-
37 1034 ing real-world communication patterns used in the distributed services of the OOI
38 1035 cyberinfrastructure [44].

39 1036 **Future Work.** Our objectives include the incorporation in the implementation of
40 1037 more elaborate handling of error cases into monitor functionality, such as halting
41 1038 all local sessions or coercing to valid actions [46, 39]. In order to reach this goal,

1
2
3
4
5 1039 we need to combine a simplification of [13] and nested sessions [20] to handle ex-
6 1040 ceptions inside MPST. We aim to construct a simple and reliable way to raise and
7 1041 catch exceptions in asynchronous networks. Another direct extension of this work
8 1042 would be the addition of states (memories) to the syntax, as described in [8, 15].
9 1043 It would require the monitors to maintain a model of the state of the applications
10 1044 being monitored, which can be easily formalised in our setting. For the sake of
11 1045 clarity, we did not add to our local type syntax other syntactical constructs such
12 1046 as parallel composition but such an extension is possible and could be consid-
13 1047 ered, as it allows one to reach greater expressiveness [22]. Our work is motivated
14 1048 by ongoing collaborations with the Savara¹ and Scribble² projects [51, 31] and
15 1049 OOI [44]. We are continuing the development of Scribble, its toolsuite and asso-
16 1050 ciated environments towards an integration into [44]. The theoretical framework
17 1051 developed in this article is extensible as a basis for other applications as demon-
18 1052 strated in our recent dynamic monitoring implementations for distributed actors
19 1053 [42] and timers [41]. For instance, the work in [41] extends run-time monitoring
20 1054 to real-time processes: monitors verify the punctuality of interactions against time
21 1055 constraints expressed as a timed extension of Scribble based on timed MPST [10].

1056 References

- 1057 [1] I. Aktug, M. Dam, and D. Gurov. Provably correct runtime monitoring. *J.*
1058 *Log. Algebr. Program.*, 78(5):304–339, 2009.
- 1059 [2] D. Ancona, S. Drossopoulou, and V. Mascaridi. Automatic generation of self-
1060 monitoring MASs from multiparty global session types in Jason. In *DALT*,
1061 volume 7784 of *LNCS*, pages 1–17. Springer, 2012.
- 1062 [3] G. Anderson and J. Rathke. Dynamic software update for message passing
1063 programs. In *APLAS*, volume 7705 of *LNCS*, pages 207–222. Springer, 2012.
- 1064 [4] L. Baresi, C. Ghezzi, and S. Guinea. Smart monitors for composed services.
1065 In *ICSOC*, pages 193–202. ACM, 2004.
- 1066 [5] L. Baresi, S. Guinea, M. Pistore, and M. Trainotti. Dynamo + astro: An inte-
1067 grated approach for BPEL monitoring. *2009 IEEE International Conference*
1068 *on Web Services*, pages 230–237, 2009.
- 1069 [6] L. Bettini, M. Coppo, L. D’Antoni, M. D. Luca, M. Dezani-Ciancaglini, and
1070 N. Yoshida. Global progress in dynamically interleaved multiparty sessions.
1071 In *CONCUR*, volume 5201 of *LNCS*, pages 418–433. Springer, 2008.

1072 ¹<http://www.jboss.org/savara>

1073 ²<http://scribble.github.io/>

- 1
2
3
4
5 1072 [7] L. Bocchi, T.-C. Chen, R. Demangeon, K. Honda, and N. Yoshida. Mon-
6 1073 itoring networks through multiparty session types. In *FMOODS/FORTE*,
7 1074 volume 7892 of *LNCS*, pages 50–65. Springer, 2013.
- 8
9
10 1075 [8] L. Bocchi, R. Demangeon, and N. Yoshida. A multiparty multi-session logic.
11 1076 In *TGC*, volume 8191 of *LNCS*, pages 97–111. Springer, 2012.
- 12
13 1077 [9] L. Bocchi, K. Honda, E. Tuosto, and N. Yoshida. A theory of design-by-
14 1078 contract for distributed multiparty interactions. In *CONCUR*, volume 6269
15 1079 of *LNCS*, pages 162–176. Springer, 2010.
- 16
17
18 1080 [10] L. Bocchi, W. Yang, and N. Yoshida. Timed multiparty session types. In
19 1081 *CONCUR*, volume 8704 of *LNCS*, pages 419–434. Springer, 2014.
- 20
21 1082 [11] M.-E. Cambronero et al. Validation and verification of web services chore-
22 1083 ographies by using timed automata. *J. Log. Algebr. Program.*, 80(1):25–49,
23 1084 2011.
- 24
25
26 1085 [12] S. Capecchi, I. Castellani, and M. Dezani-Ciancaglini. Information flow
27 1086 safety in multiparty sessions. In *EXPRESS*, volume 64 of *EPTCS*, pages
28 1087 16–30, 2011.
- 29
30
31 1088 [13] S. Capecchi, E. Giachino, and N. Yoshida. Global escape in multiparty ses-
32 1089 sion. In *FSTTCS*, volume 8 of *LIPICS*, pages 338–351. Schloss Dagstuhl -
33 1090 Leibniz-Zentrum fuer Informatik, 2010.
- 34
35
36 1091 [14] F. Chen and G. Rosu. MOP: an efficient and generic runtime verification
37 1092 framework. In *OOPSLA*, pages 569–588. ACM, 2007.
- 38
39 1093 [15] T. Chen and K. Honda. Specifying stateful asynchronous properties for dis-
40 1094 tributed programs. In *CONCUR*, volume 7454 of *LNCS*, pages 209–224.
41 1095 Springer, 2012.
- 42
43 1096 [16] T.-C. Chen. *Theories for Session-based Governance for Large-Scale Dis-*
44 1097 *tributed Systems*. PhD thesis, Queen Mary, University of London, 2013.
- 45
46
47 1098 [17] T.-C. Chen, L. Bocchi, P.-M. Deniélou, K. Honda, and N. Yoshida. Asyn-
48 1099 chronous distributed monitoring for multiparty session enforcement. In
49 1100 *TGC*, volume 7173 of *LNCS*, pages 25–45. Springer, 2011.
- 50
51
52 1101 [18] M. Coppo, M. Dezani-Ciancaglini, and B. Venneri. Self-adaptive monitors
53 1102 for multiparty sessions. In *PDP*, pages 688–696. IEEE, 2014.
- 54
55 1103 [19] R. De Nicola, G. Ferrari, and R. Pugliese. Klaim: a kernel language for
56 1104 agents interaction and mobility. *IEEE Trans. Softw. Eng.*, 24:315–330, 1998.

- 1
2
3
4
5 1105 [20] R. Demangeon and K. Honda. Nested protocols in session types. In *CON-*
6 1106 *CUR*, volume 7454 of *LNCS*, pages 272–286. Springer, 2012.
- 7
8 1107 [21] R. Demangeon, K. Honda, R. Hu, R. Neykova, and N. Yoshida. Practical
9 1108 interruptible conversations: Distributed dynamic verification with multiparty
10 1109 session types and python. *FMSD*, pages 1–29, 2015.
- 11
12
13 1110 [22] R. Demangeon and N. Yoshida. On the expressiveness of multiparty ses-
14 1111 sions. In *FSTTCS*, volume 45 of *LIPICs*, pages 560–574. Schloss Dagstuhl -
15 1112 Leibniz-Zentrum fuer Informatik, 2015.
- 16
17
18 1113 [23] P.-M. Deniélou and N. Yoshida. Dynamic multirole session types. In *POPL*,
19 1114 pages 435–446. ACM, 2011.
- 20
21 1115 [24] P.-M. Deniélou and N. Yoshida. Multiparty session types meet communicat-
22 1116 ing automata. In *ESOP*, volume 7211 of *LNCS*, pages 194–213. Springer,
23 1117 2012.
- 24
25
26 1118 [25] G. L. Ferrari, E. Moggi, and R. Pugliese. Guardians for ambient-based mon-
27 1119 itoring. *Electr. Notes Theor. Comput. Sci.*, 66(3):52–75, 2002.
- 28
29 1120 [26] Y. Gan, M. Chechik, S. Nejati, J. Bennett, B. O’Farrell, and J. Waterhouse.
30 1121 Runtime monitoring of web service conversations. In *CASCON*, pages 42–
31 1122 57. ACM, 2007.
- 32
33
34 1123 [27] C. Ghezzi and S. Guinea. Run-time monitoring in service-oriented archi-
35 1124 tectures. In *Test and Analysis of Web Services*, pages 237–264. Springer,
36 1125 2007.
- 37
38
39 1126 [28] K. W. Hamlen and M. Jones. Aspect-oriented in-lined reference monitors.
40 1127 In *PLAS*, pages 11–20. ACM, 2008.
- 41
42 1128 [29] K. Havelund and A. Goldberg. Verify your runs. In *Verified Software: The-*
43 1129 *ories, Tools, Experiments*, pages 374–383. Springer, 2008.
- 44
45 1130 [30] K. Honda, R. Hu, R. Neykova, T.-C. Chen, R. Demangeon, P.-M. Denilou,
46 1131 and N. Yoshida. Structuring communication with session types. In *COB*,
47 1132 volume 8665 of *LNCS*, pages 105–127. Springer, 2012.
- 48
49
50 1133 [31] K. Honda, A. Mukhamedov, G. Brown, T.-C. Chen, and N. Yoshida. Scrib-
51 1134 bling interactions with a formal foundation. In *ICDCIT*, volume 6536 of
52 1135 *LNCS*, pages 55–75. Springer, 2011.
- 53
54
55 1136 [32] K. Honda and N. Yoshida. On reduction-based process semantics. *Theor.*
56 1137 *Comput. Sci.*, 151(2):437–486, 1995.

- 1
2
3
4
5 1138 [33] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session
6 1139 types. In *POPL*, pages 273–284. ACM, 2008.
- 7
8
9 1140 [34] R. Hu, R. Neykova, N. Yoshida, and R. Demangeon. Practical interruptible
10 1141 conversations: distributed dynamic verification with session types and python.
11 1142 In *RV*, volume 8174 of *LNCS*, pages 148–130. Springer, 2013.
- 12
13 1143 [35] L. Jia, H. Gommerstadt, and F. Pfenning. Monitors and blame assignment
14 1144 for higher-order session types. In *POPL*, pages 582–594. ACM, 2016.
- 15
16 1145 [36] I. H. Krüger, M. Meisinger, and M. Menarini. Runtime verification of inter-
17 1146 actions: from MSCs to aspects. In *RV*, volume 4839 of *LNCS*, pages 63–74.
18 1147 Springer, 2007.
- 19
20
21 1148 [37] I. H. Krüger, M. Meisinger, and M. Menarini. Interaction-based runtime
22 1149 verification for systems of systems integration. *J. Log. Comput.*, 20(3):725–
23 1150 742, 2010.
- 24
25
26 1151 [38] M. Leucker and C. Schallhart. A brief account of runtime verification. *J.*
27 1152 *Log. Algebr. Program.*, 78(5):293–303, 2009.
- 28
29
30 1153 [39] J. Ligatti, L. Bauer, and D. Walker. Run-time enforcement of nonsafety
31 1154 policies. *ACM Trans. Inf. Syst. Secur.*, 12:19:1–19:41, 2009.
- 32
33 1155 [40] V. Mascardi and D. Ancona. Attribute global types for dynamic check-
34 1156 ing of protocols in logic-based multiagent systems. *TPLP*, 13(4-5-Online-
35 1157 Supplement), 2013.
- 36
37
38 1158 [41] R. Neykova, L. Bocchi, and N. Yoshida. Timed runtime monitoring for
39 1159 multiparty conversations. In *BEAT*, volume 162 of *EPTCS*, pages 19–26,
40 1160 2014.
- 41
42 1161 [42] R. Neykova and N. Yoshida. Multiparty session actors. In *COORDINATION*,
43 1162 volume 8459 of *LNCS*, pages 131–146. Springer, 2014. A full version in
44 1163 *LMCS*.
- 45
46
47 1164 [43] R. Neykova, N. Yoshida, and R. Hu. SPY: Local Verification of Global
48 1165 Protocols. In *RV*, volume 8174 of *LNCS*, pages 363–358. Springer, 2013.
- 49
50 1166 [44] OOI. <http://www.oceanobservatories.org/>.
- 51
52
53 1167 [45] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- 54
55 1168 [46] F. B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*,
56 1169 3:30–50, 2000.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1170 [47] Scribble Project homepage. <http://www.scribble.org>.

1171 [48] J. Simmonds, Y. Gan, M. Chechik, S. Nejati, B. O’Farrell, E. Litani, and
1172 J. Waterhouse. Runtime monitoring of web service conversations. *IEEE T.*
1173 *Services Computing*, 2(3):223–244, 2009.

1174 [49] M. Sridhar and K. W. Hamlen. Model-checking in-lined reference monitors.
1175 In *VMCAI*, volume 5944 of *LNCS*, pages 312–327. Springer, 2010.

1176 [50] W. M. P. van der Aalst, M. Dumas, C. Ouyang, A. Rozinat, and E. Verbeek.
1177 Conformance checking of service behavior. *ACM Trans. Internet Techn.*,
1178 8(3), 2008.

1179 [51] N. Yoshida, R. Hu, R. Neykova, and N. Ng. The scribble protocol language.
1180 In *TGC*, volume 8358 of *LNCS*, pages 22–41. Springer, 2013.