

Nested Protocols in Session Types

Romain Demangeon and Kohei Honda

Queen Mary, University of London

Abstract. We propose an improvement to session-types, introducing nested protocols, the possibility to call a subprotocol from a parent protocol. This feature adds expressiveness and modularity to the existing session-type theory, allowing arguments to be passed and enabling higher-order protocols definition. Our theory is introduced through a new type system for protocols handling subprotocol calls, and its implementation in a session-calculus. We propose validation and satisfaction relations between specification and implementation. Sound behaviour is enforced thanks to the usage of kinds and well-formedness, allowing us to ensure progress and subject reduction. In addition, we describe an extension of our framework allowing subprotocols to send back results.

1 Introduction

Decentralised computation is becoming more and more popular thanks to the fast growth of web services and other distributed computing technologies. In such a distributed framework, agents (users, servers, applications) are interacting through message-passing communications, without central control. The programmatic coordination of a large number of independent entities interacting with each other inside a network is a challenging task: without global control, the only place where coordination can come from is local endpoints. How can we specify and ensure correctly coordinated behaviour without having any global control? Session types [11] provide a powerful expressive framework to help solving this issue, focusing on the notion of *session* seen as a unit of conversation among participants called *roles*. The expected scenario of the session is described in a global protocol given as a *global type*, projected into end-point specifications called *local types*, describing the behaviour of each role. Those are enforced locally, either through a static analysis of programs (static validation [8, 3, 13]) or at run-time (monitoring [2]). If each agent in the network conforms to its local type, it is guaranteed that their overall interactions conform to the global specification. In the past few years, the theory of session types has been extended in several directions. On the one hand, new features added to the language of the global types allow one to specify more accurately the interactions inside a protocol, for instance by including logical assertions [3], or information flow [4]. This “internal expressiveness” ensure the satisfaction of auxiliary properties: security (“the messages between Alice and Bob cannot be read by Carol”), or governance (“Alice can send a buying request to Bob only if she has enough money on her

bank account”). On the other hand, extensions of the session mechanisms allow greater control over, for example, how participants join or leave a session (through dynamic multiparty session [9] or through a reputation system [10]) or how sessions can be parametrised, increasing “external expressiveness”.

Real-world specifications for decentralised networks are large, complex and often highly modular: for example, such specifications are found in many use cases from the development of a large-scale distributed infrastructure for ocean sciences, Ocean Observatories Initiative [15], with whom we are collaborating. Among the use cases in the project, several protocols, used in different contexts, share the same shape. Moreover, some protocols call other protocols. In order to be able to specify, verify, simplify and organise such complex protocol frameworks effectively, solid improvements to the theory of session-types are needed.

In this paper, we present a novel approach to session-types that addresses the structuring principle itself of protocols, increasing both internal and external expressiveness. We introduce *nesting* of protocols, that is, the possibility to define a subprotocol independently of its parent protocol, which calls the subprotocol explicitly. Through a call, arguments can be passed, such as values, roles and other protocols, allowing higher-order description. At the programming level, subprotocols are realised as *subsessions*: one agent creates a new private session, inviting roles of the parent session (*internal invitations*) as well as other agents from the network (*external invitations*). Uninvited participants of the parent session do not have access to the subsession, allowing one to model private interactions inside public sessions. This contrasts with the current use of subprotocols in the protocol description language Scribble [16], where they only correspond to the in-lining mechanism. As an example, we noticed that in several use cases described in [15], a negotiation procedure *Nego* between two agents is invoked inside a main protocol, and other participants do not take part in that negotiation. The *Nego* procedure has its own description, subject to independent modifications and can be invoked in different contexts for different purposes. The theory we propose introduces such modularity in the framework of session types, yielding a solid, formal verification method for distributed programs.

One strong motivation for introducing subprotocols and subsession is that they are a powerful structuring tool: cross-cutting features such as login, negotiations or security controls can be abstracted from the targeted protocols in a compositional way, becoming subprotocols. If we update a login protocol to enforce stronger security checks, specifications of applications using it do not need be updated. Moreover, nesting allows one to call multiple copies of the same protocol with different arguments, improving flexibility and readability. This allows us to clean up and reorganise a large protocol database in [15], by unifying many protocols with the same shape into one parametrised protocol. Another direct benefit from nesting is allowing a better separation of the different branches by inviting participants only when necessary, reducing complexity and resource usage. For instance, in protocol \mathcal{P} involving Alice, Bob and Carol, if Carol interacts only if a certain condition is met, we can have Carol act in a separate subprotocol, inviting Carol only if her presence is required. In our framework, internal

invitations to a subsession are sent *inside the parent-session*, targeting a specific participant through a linear channel. This extends the existing session-calculi where invitations are always done externally, through shared channels.

We propose in Section 2 a syntax for nested protocols, with dedicated constructors for protocol definitions and protocol calls. In order to ensure sound composition, we introduce the notion of *kinds*, “types for types”, and define a notion of well-formedness. In Section 3, we describe a session-calculus (based on the π -calculus [14]) handling subsessions. We describe in Section 4 a static validation and a run-time satisfaction, each linking specifications and processes in the session-calculus. We sketch the main properties of typed processes. Finally we propose an extension to our theory in Section 5, allowing a subsession to have a goal: a *result* that is returned to its parent session.

2 Nested Protocols

Global types Throughout the paper, we use G for global types, T for local types, l for communication labels, s, k for session names, a, b for shared channels, ℓ for transition labels, \mathbf{r}, \mathbf{r}' for role identifiers and \mathcal{P} for protocol identifiers. We use v to describe *values*, which can be base type variables (integers, strings, \dots), labels or protocol identifiers. x, y are variables, possibly abstracting any value. For any identifier e , we use \tilde{e} to abstract the sequence e_1, \dots, e_n of unspecified length n . We use \mathcal{R}^+ (resp. \mathcal{R}^*) for the reflexive (resp. reflexive-transitive) closure of the relation \mathcal{R} . We assume a Barendregt convention for bound variables.

Global types describe protocols from the network point of view: they consist of sequences of interactions between *roles*. We choose the already existing syntax of multi-party session types (e.g., in [3, 1]) as a basis. The syntax for our global types is given by the following grammar:

$$\begin{array}{lcl}
 G & ::= & \text{let } \mathcal{P} = \lambda \tilde{\mathbf{r}}^1, \tilde{y} \mapsto \text{new } \tilde{\mathbf{r}}^2. G \text{ in } G' & \text{(declaration)} \\
 & | & \mathbf{r} \text{ calls } \mathcal{P}(\tilde{x}, \tilde{y}). G & \text{(call)} \\
 & | & \mathbf{r}_1 \rightarrow \mathbf{r}_2 : \Sigma_{i \in I} \{l_i(x_i). G_i\} & \text{(com)} \\
 & | & G_1 \oplus^x G_2 \mid G_1 \parallel G_2 \mid \mu \mathbf{t}. G \mid \mathbf{t} & \text{(choice, par, rec, rec-var)}
 \end{array}$$

Communications between two roles are specified with $\mathbf{r}_1 \rightarrow \mathbf{r}_2 : \Sigma_{i \in I} \{l_i(x_i : S_i). G_i\}$, stating that \mathbf{r}_1 has a *directed choice* between several labels \tilde{l} proposed by \mathbf{r}_2 . Each branch expects a value x_i and executes the continuation G_i . When I is a singleton, we write $\mathbf{r}_1 \rightarrow \mathbf{r}_2 : l(x : S)$. Primitive \oplus^x is *located choice*: the choice for one participant \mathbf{r} between two distinct protocol branches. Parallel composition is denoted by \parallel and recursion by the two operators $\mu \mathbf{t}$ and \mathbf{t} . We assume a congruence over global types, handling implicit unfolding of recursion.

The new primitives addressing protocol stratification are **let** and **calls**. We describe the declaration of an auxiliary protocol, to be called by a main protocol, by the notation **let** $\mathcal{P} = \lambda \tilde{\mathbf{r}}^1, \tilde{y} \mapsto \text{new } \tilde{\mathbf{r}}^2. G_1 \text{ in } G_2$. In this notation, the protocol G_1 is identified by \mathcal{P} in the main protocol G_2 . The participants of G_1 are explicitly separated into two groups, $\tilde{\mathbf{r}}^1$ are internally invited from the parent session and thus given as arguments to \mathcal{P} together with values \tilde{y} ; whereas $\tilde{\mathbf{r}}^2$ are

externally invited from the network at the beginning of G_1 . The counterpart of this constructor is the *protocol call* \mathbf{r} calls $\mathcal{P}(\tilde{\mathbf{r}}, \tilde{v})$ stating that participant \mathbf{r} executes an auxiliary protocol \mathcal{P} with role arguments $\tilde{\mathbf{r}}$, value arguments \tilde{v} . Note that \tilde{v} can contain protocol identifiers, thus allowing higher-order programming.

Kinds As protocols can be abstracted, called and used as arguments, we introduce a simple and concise discipline for protocols, which ensures that they are used in an adequate way, *well-formedness*. In order to formalise this notion, we type all objects appearing in specifications with *kinds* (types for types) $K, S ::= \mathbf{Role} \mid \mathbf{Val} \mid \diamond \mid (K_1 \times \cdots \times K_n) \rightarrow K$. We use \mathbf{Val} to denote the value-kinds, which are first-order types for values (like \mathbf{Nat} for integers) or data types (such as \mathbf{Req} in Section 2), \diamond to denote protocol type and \rightarrow to denote parametrisation. The presence of higher-order calls allows us to treat protocols whose kinds have shapes like $\mathbf{Role} \times (\mathbf{Role} \rightarrow \diamond) \rightarrow \diamond$, describing a protocol parametrised by a role and another protocol, the latter parametrised by a role. In the following, we will sometimes adopt an *à la Church* notation for protocol constructors, as in $\mathbf{let} P = \lambda \text{buyer} : \mathbf{Role}, \text{price} : \mathbf{Nat} \mapsto \mathbf{new} \tilde{\mathbf{r}}.G \text{ in } G'$, in order to specify the kinds of the arguments passed to a protocol.

We define well-formedness to rule out unsound protocols. For instance, a protocol where \mathcal{P} has kind $\mathbf{Role} \rightarrow \diamond$ but is used with kind $\mathbf{Nat} \rightarrow \diamond$ is not *well-kinded*. A protocol containing $(\mathbf{r}_1 \rightarrow \mathbf{r}_2 : \text{ok}) \oplus^{\mathbf{r}_0} (\mathbf{r}_2 \rightarrow \mathbf{r}_1 : \text{ko})$ is not *projectable* as \mathbf{r}_1 has no mean to know which branch \mathbf{r}_0 chooses, and thus is not able to know if it must perform an input or an output. In order to define projectability, which ensures that a global type can be coherently projected into local types, we define the *restriction* of a protocol to a role, noted $(G)|_{\mathbf{r}}$ as the global type obtained by removing every constructor of G where \mathbf{r} does not appear. A protocol is projectable if for every choice (directed or located), the difference between the branches are only visible to the roles involved in that choice.

Definition 1 (Well-Formedness)

A global-type G is well-kinded if there exists τ from all identifiers of G to types satisfying, for all subprotocols of G :

1. $\mathbf{let} P = (\lambda \tilde{\mathbf{y}} \mapsto \mathbf{new} \tilde{\mathbf{r}} \cdot) \text{ in } \cdot$ is s.t. $\tau(P) = \tau(\tilde{\mathbf{y}}) \rightarrow \diamond$, and $\forall i, \tau(r_i) = \mathbf{Role}$.
2. \mathbf{r} calls $\mathcal{P}(\tilde{\mathbf{y}}).G$ is s.t. $\tau(\mathbf{r}) = \mathbf{Role}$, $\tau(P) = \tau(\tilde{\mathbf{y}}) \rightarrow \diamond$.
3. for all identifiers \mathbf{r} in $\mathbf{r} \rightarrow \mathbf{r}' : l(x : S).G$ and $G_1 \oplus^{\mathbf{r}} G_2$, $\tau(\mathbf{r}) = \mathbf{Role}$.

A global type G is projectable if:

1. for each subterm of G of the form $G_1 \oplus^{\mathbf{r}_0} G_2$, for any free $\mathbf{r} \neq \mathbf{r}_0$, $(G_1)|_{\mathbf{r}} = (G_2)|_{\mathbf{r}}$.
2. for each subterm of G of the form $\mathbf{r} \rightarrow \mathbf{r}' : \Sigma_{i \in I} \{l_i(x_i : S_i).G_i\}$, for any role free role $\mathbf{r} \notin \{\mathbf{r}, \mathbf{r}'\}$ and for all $\{i, j\} \subseteq I$, $(G_i)|_{\mathbf{r}} = (G_j)|_{\mathbf{r}}$.

A protocol G is well-formed when it is well-kinded and projectable, and satisfies the standard linearity condition [1].

There exists in [9, 12] *mergeability conditions* that allows the authors to be less restrictive in the definition of projectability. Our framework could accommodate this refinement. We do not present it here, for the sake of clarity.

Motivating Examples In this section, we motivate our contribution with three examples extracted from concrete specifications and illustrate higher-order programming with a fourth one.

Resource usage The following example is inspired by the use cases (UC R2.34, UC R2.32) from the OOI project [15]. A negotiation procedure *Nego* is first defined independently, to be used in several different protocols. This negotiation procedure involves two participants trying to agree on a contract: first participant specifies a request, second participant offers a corresponding contract, then both participants enter a loop when the first one can either accept the contract, which ends the protocol or make a counter-offer.

```

let Nego = λr1, r2 ↦
    r1 → r2 : ask(terms).
    μt.
    r2 → r1 : proposition(contract2).
    r1 → r2 : {accept.end
                counter(contract1).t}
in
client → agent : request(coord). agent → instr : connect
instr → agent : available. agent → client : ack.
agent calls Nego(agent, client).
μt.
client → instr : {abort(coord).end
                  command(code).
                  instr → client : result(data).t}

```

The main protocol *UseRes* consists of several interactions between three participants (client, agent, instr), processed in the following order: first client sends a request to agent for an instrument he wants to use, agent tries to connect to instr which acknowledges when available. Then, agent negotiates a contract with client (by calling protocol *Nego*). After a successful negotiation, client and instr interact inside a loop, the client sending commands and receiving data. The negotiation phase is considered external: should the auxiliary protocol be modified, for instance to enforce another negotiation policy, the main protocol would remain the same.

Client-Middleware-Server The protocol *CMS*, presented of the left side of Figure 1 describing a typical service interaction. This protocol initially involves two participants, client starts the interactions by sending a request to the middleware middle. If the latter is able to treat the request directly, it answers to client, if not, it contacts server, calling subprotocol *Contact* with itself as role argument. In the subprotocol, middle performs an external invitation of server, forwards the request and waits for an answer. After the subprotocol is completed, the answer is forwarded to the client. Nesting, in this example, allows us to invite server to participate only when necessary: if middle can treat the request, server is not even invited. Using subprotocols in such a way allows us to cut a great deal of unnecessary traffic caused by external invitations, saving bandwidth.

Dynamic distribution We then describe on the right side of Figure 1 a third example, inspired by a concrete protocol from the Array Network Facility, used

```

let Contact = λagent, req ↦
  new server.
  agent → server : request(req).
  server → agent : answer(ans).
end
in
client → middle : request(req0).
⊕middle (middle → client : answer(ans0).end)
(middle calls Contact(middle, req0).
middle → client : answer(ans0).end) end

let Treat =
λr1, r2 ↦
  new worker.
  r1 → worker : raw(data).
  worker → r2 : processed(data).
end
in
paralleln(
  source calls Treat(source, target)
).
end

```

Fig. 1. Protocols *CMS* and *ANF*

for processing seismic data. Here, the operator $\text{parallel}_n(G)$ is used as a shortcut for n parallel copies of the protocol G . In this protocol, data comes in a raw state from a participant source and should reach participant target processed. In the body of each of the n parallel executions, source calls the subprotocol *Treat* inviting a new participant worker and using it to process data. This protocol is run in networks where many computing units can accept temporarily the worker role. In this example, stratification is used to present in a clean way the execution of thousands of copies of the same protocol. As each copy is implemented by a different session, the different calls to the subprotocol are actually independent from each other.

Marketplace Finally, we propose a protocol for a virtual marketplace in which participants have the possibility to engage in trade actions with other participants. General protocols *Buy* and *Sell* are defined to handle these buying and selling. The encounter between two agents follows the same procedure (handshake, authentication, possibility to cancel the transaction) whatever the reason of their meeting is. This common procedure is abstracted in *Meet* and a protocol identifier *Action* is given as an argument to *Meet* calls, meant to be substituted by *Buy* or *Sell* (or any similar protocol). Thus, *Meet* is an higher-order protocol, parametrised with protocol *Action*.

```

let Buy = λ agent : Role, seller : Role, item : Tradable ↦ ...
in let Sell = λ agent : Role, buyer : Role, item : Tradable ↦ ...
in let Meet = λ agent : Role, partner : Role,
  item : Tradable, Action : (Role → Role → Tradable → ◇) ↦ ...
  agent calls Action(partner, item) ...
  in ...
  alice calls Meet(bob, kettle, Buy). carol calls Meet(bob, teacup, Sell) ...

```

The protocols presented in this section are well-formed: notice that protocol *CMS* is projectable, in each branch of the choice \oplus^{middle} , the restriction on client is “middle → client : answer.end”. Kinds for subprotocols presented in the examples are as follows: $\text{Nego} : \text{Role} \times \text{Role} \rightarrow \diamond$, $\text{Contact} : \text{Role} \times \text{Req} \rightarrow \diamond$, $\text{Treat} : \text{Role} \times \text{Role} \rightarrow \diamond$, $\text{Buy}, \text{Sell} : \text{Role} \times \text{Role} \times \text{Tradable} \rightarrow \diamond$, $\text{Meet} : \text{Role} \times \text{Role} \times \text{Tradable} \times (\text{Role} \times \text{Role} \times \text{Tradable} \rightarrow \diamond) \rightarrow \diamond$

Local types and Projection Local types describe a global conversation from the partial point-of-view of a participant and are used to validate and monitor distributed programs. Their syntax is given by:

$$\begin{aligned}
T ::= & \text{get}[\mathbf{r}]?_{i \in I} \{l_i(x_i : S_i).T_i\} \mid \text{send}[\mathbf{r}]!_{i \in I} \{l_i(x_i : S_i).T_i\} \\
& \mid T \parallel T \mid T \oplus T \mid \mu \mathbf{t}.T \mid \mathbf{t} \mid \text{end} \\
& \mid \text{call } \mathcal{P} : G \text{ with } (\tilde{v} \text{ as } \tilde{y} : \tilde{S}) \& (\tilde{\mathbf{r}}^2).T \\
& \mid \text{ent } \mathcal{P}[\mathbf{r}] \langle \tilde{v} \rangle \text{ from } \mathbf{r}.T \mid \text{req } \mathcal{P}[\mathbf{r}] \langle \tilde{v} \rangle \text{ to } \mathbf{r}.T
\end{aligned}$$

Creating a subsession for protocol \mathcal{P} having global type G is specified by $\text{call } \mathcal{P} : G \text{ with } (\tilde{v} \text{ as } \tilde{y} : \tilde{S}) \& (\tilde{\mathbf{r}}^2)$, with \tilde{v} as value arguments and involving external invitations for roles $\tilde{\mathbf{r}}^2$. Internal invitations are handled using two specific constructors, as they are meant to be performed on the parent session channel: ent specifies the act of accepting such an invitation, req specifies the dual action. Syntax contains endpoint primitives for communications, specified by get for the receiver side and send for the sender side, as well as constructors for parallel, choice and recursion. We handle equivalence of types through recursions and parallel compositions implicitly. In the following, we omit trailing occurrences of end .

Projection from global to local types is defined w.r.t. a protocol environment, associating protocols identifiers to their contents. Environment is updated by let in constructors. We present below the projection rule for call and let . For the former the result of the projection depends on the participant we project on, \mathbf{r}_{proj} . If it is the subprotocol initiator \mathbf{r}^A it is responsible for creating the subsession (call) and sending the internal invitations (req). If it participates in the subprotocol, it has to accept an internal invitation (ent). Projection on other constructors is standard.

$$\begin{aligned}
& (\text{let } \mathcal{P} = \tilde{\mathbf{r}}^1.\tilde{y} \mapsto \text{new } \tilde{\mathbf{r}}^2.G_{\mathcal{P}} \text{ in } G') \Downarrow_{\mathbf{r}^p}^{\text{Env}} = G' \Downarrow_{\mathbf{r}^p}^{\text{Env}, \mathcal{P} \mapsto (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})} \\
& \left(\mathbf{r}^A \text{ calls } \mathcal{P}(\tilde{\mathbf{r}}^0, \tilde{v}).G \right) \Downarrow_{\mathbf{r}^p}^{\text{Env}, \mathcal{P} \mapsto (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})} = \\
& \left\{ \begin{array}{l}
\text{if } \mathbf{r}^p = \mathbf{r}^A, \mathbf{r}^A \notin \tilde{\mathbf{r}}^0 \\
\quad \text{call } \mathcal{P} : G_{\mathcal{P}} \text{ with } (\tilde{v} \text{ as } \tilde{y}) \& (\tilde{\mathbf{r}}^2).[(G) \Downarrow_{\mathbf{r}^p}^{\text{Env}, \mathcal{P} \mapsto (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})} \\
\quad \parallel \text{req } \mathcal{P}[\mathbf{r}_0^1] \langle \tilde{v} \rangle \text{ to } \mathbf{r}_0^0 \parallel \dots \parallel \text{req } \mathcal{P}[\mathbf{r}_n^1] \langle \tilde{v} \rangle \text{ to } \mathbf{r}_n^0] \\
\text{if } \mathbf{r}^p = \mathbf{r}^A \text{ and } \mathbf{r}^A = \mathbf{r}_i^0 \\
\quad \text{call } \mathcal{P} : G_{\mathcal{P}} \text{ with } (\tilde{v} \text{ as } \tilde{y}) \& (\tilde{\mathbf{r}}^2).[(G) \Downarrow_{\mathbf{r}^p}^{\text{Env}, \mathcal{P} \mapsto (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})} \\
\quad \parallel \text{ent } \mathcal{P}[\mathbf{r}_i^1] \langle \tilde{v} \rangle \text{ from } \mathbf{r}^A \parallel \text{req } \mathcal{P}[\mathbf{r}_0^1] \langle \tilde{v} \rangle \text{ to } \mathbf{r}_0^0 \parallel \dots \parallel \text{req } \mathcal{P}[\mathbf{r}_n^1] \langle \tilde{v} \rangle \text{ to } \mathbf{r}_n^0] \\
\text{if } \mathbf{r}^p \neq \mathbf{r}^A \text{ and } \mathbf{r}^p = \mathbf{r}_i^0 \\
\quad \text{ent } \mathcal{P}[\mathbf{r}_i^1] \langle \tilde{v} \rangle \text{ from } \mathbf{r}^A.(G) \Downarrow_{\mathbf{r}^p}^{\text{Env}, \mathcal{P} \mapsto (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})} \\
\text{Otherwise} \\
\quad (G) \Downarrow_{\mathbf{r}^p}^{\text{Env}, \mathcal{P} \mapsto (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})}
\end{array} \right.
\end{aligned}$$

If the initiator \mathbf{r}^A of the subsession also takes part in it, the projection on \mathbf{r}^A specifies that it invites itself. It is easy to add to our language a dedicated constructor handling session-invitation directly, without inducing communication at the network level. For the sake of clarity, we do not include such a constructor in this paper.

We present below the projection of CMS on its two roles. G_{CMS} is the global type of the whole protocol and G_C the global type of $Contact$.

$$\begin{aligned}
G_{CMS} \Downarrow_{\text{client}}^{\emptyset} &= \text{send}[\text{middle}]!\{\text{request}(req_0)\}.\text{get}[\text{middle}]?\{\text{answer}(ans_0)\} \\
G_{CMS} \Downarrow_{\text{middlew}}^{\emptyset} &= \text{get}[\text{client}]?\{\text{request}(req_0)\}. \\
&\quad \text{send}[\text{client}]!\{\text{answer}(ans_0)\} \\
&\quad \oplus \\
&\quad (\text{call } Contact : G_C \text{ with } (req_0 \text{ as } req : \text{Req})\&(\text{server}). \\
&\quad \quad (\text{req } Contact[\text{agent}]\langle req \rangle \text{ to middle } \parallel \\
&\quad \quad \text{ent } Contact[\text{agent}]\langle req \rangle \text{ from middle } \parallel \\
&\quad \quad \text{send}[\text{client}]!\{\text{answer}(ans_0)\}))
\end{aligned}$$

3 Session-Calculus

Our session-calculus, based on the π -calculus [14], contains usual primitives from existing session-calculi [3], as well as dedicated primitives for session creation and internal (on-session) invitations. Names are divided into shared channels a, b, u (standard π -names) and session channels s, k . The former are used to send and receive external invitations, the latter to handle all session interactions.

$$\begin{aligned}
P ::= & \mathbf{0} \mid P \mid P \mid a(x).P \mid \bar{a}(s).P \mid P + P \\
& \mid k?[\mathbf{r}, \mathbf{r}]_{i \in I} \{l_i(x_i).P_i\} \mid k![\mathbf{r}, \mathbf{r}]l\langle v \rangle.P \mid (\nu u) P \\
& \mid \text{new } s \text{ on } s \text{ with } (\tilde{v})\&(\tilde{a} \text{ as } \tilde{\mathbf{r}}).P \\
& \mid s \downarrow [\mathbf{r}, \mathbf{r} : \mathbf{r}](x).P \mid s \uparrow [\mathbf{r}, \mathbf{r} : \mathbf{r}]\langle s \rangle.P \mid \mu X(x).P\langle v \rangle \mid X\langle v \rangle
\end{aligned}$$

We denote by $k?[\mathbf{r}_1, \mathbf{r}_2]_{i \in I} \{l_i(x_i).P_i\}$ a branching input on session k from \mathbf{r}_1 to \mathbf{r}_2 , with continuations $(P_i)_{i \in I}$. The dual primitive is $k![\mathbf{r}_1, \mathbf{r}_2]l\langle v \rangle.P$. Creation of a subsession is done with **new** s on k with $(\tilde{v})\&(\tilde{a} \text{ as } \tilde{\mathbf{r}}^2)$ with s being the subsession, k the parent session, \tilde{v} the arguments and \tilde{a} the channels on which the external invitations are sent. Operator $s \downarrow [\mathbf{r}^1, \mathbf{r}^2 : \mathbf{r}^3](x).P$ is the action of waiting on s for an internal invitation sent by \mathbf{r}^1 to \mathbf{r}^2 in order to play role \mathbf{r}^3 in a subsession x . Finally, $s \uparrow [\mathbf{r}^1, \mathbf{r}^2 : \mathbf{r}^3]\langle s \rangle.P$ is its dual action. Inputs and outputs on shared channels, choice, parallel composition and inactive process $\mathbf{0}$ are inherited from the π -calculus. We omit trailing occurrences of $\mathbf{0}$. Structural congruence \equiv for processes is defined in the usual way.

Semantics is given by reduction rules below, defined w.r.t. a notion of evaluation context $\mathbf{E} ::= [] \mid P \mid \mathbf{E} \mid (\nu a) \mathbf{E}$. The crucial rule of our system is (**subs**) where a session creation operator **new** is destructed in order to create external invitations on shared channels. (**join**) handles internal session invitation, other rules are standard:

$$\begin{aligned}
(\text{comS}) \quad & \overline{\mathbf{E}[s[\mathbf{r}_1, \mathbf{r}_2]l_j\langle \tilde{v} \rangle.P \mid s?[\mathbf{r}_1, \mathbf{r}_2]_{i \in I} \{l_i(\tilde{x}_i).P_i\}] \rightarrow \mathbf{E}[P \mid P_j\{\tilde{v}/\tilde{x}_j\}]} \\
(\text{comC}) \quad & \overline{\mathbf{E}[\bar{a}(\tilde{v}).P \mid a(\tilde{y}).Q] \rightarrow \mathbf{E}[P \mid Q\{\tilde{v}/\tilde{y}\}]} \\
(\text{subs}) \quad & \overline{\tilde{\mathbf{r}}^2 = (\mathbf{r}_1^2, \dots, \mathbf{r}_n^2) \quad \tilde{a} = (a_1, \dots, a_n)} \\
& \overline{\mathbf{E}[\text{new } s \text{ on } k \text{ with } (\tilde{v})\&(\tilde{a} \text{ as } \tilde{\mathbf{r}}^2).P] \rightarrow \mathbf{E}[P \mid \bar{a}_1\langle s[\mathbf{r}_1^2] \rangle \mid \dots \mid \bar{a}_n\langle s[\mathbf{r}_n^2] \rangle]} \\
(\text{join}) \quad & \overline{\mathbf{E}[s \uparrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}']\langle k \rangle.P \mid s \downarrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}']\langle x \rangle.Q] \rightarrow \mathbf{E}[P \mid Q\{k/x\}]} \\
(\text{choice}) \quad & \overline{P_i \rightarrow P'_i} \\
& \overline{\mathbf{E}[(P_1 + P_2)] \rightarrow \mathbf{E}[P'_i]}
\end{aligned}$$

As an example consider the following processes:

$$\begin{aligned}
P_{\text{alice}} &= a(x).x![\text{client}, \text{middle}]\text{request}(\text{"kettle"}).x?[\text{middle}, \text{client}]\text{answer}(ans_0) \\
P_{\text{bob}} &= \bar{a}(s).s?[\text{client}, \text{middle}]\text{request}(req_0). \\
&\quad (s![\text{middle}, \text{client}]\text{answer}(ans_0) \\
&\quad + (\text{new } k \text{ on } s \text{ with } (req_0) \& (c \text{ as server}). \\
&\quad \quad s \uparrow [\text{middle}, \text{middle} : \text{agent}](k) \mid s \downarrow [\text{middle}, \text{middle} : \text{agent}](z). \\
&\quad \quad z![\text{agent}, \text{server}]\text{request}(req_0).z?[\text{server}, \text{agent}]\text{answer}(ans_r). \\
&\quad \quad s![\text{middle}, \text{client}]\text{answer}(ans_r)) \\
P_{\text{carol}} &= c(y).y?[\text{agent}, \text{server}]\text{request}(req).y![\text{server}, \text{agent}]\text{answer}(ans)
\end{aligned}$$

P_{alice} , P_{bob} and P_{carol} are processes ready to play, respectively roles client, middle and server in the *CMS* protocol. P_{alice} (resp. P_{carol}) is a simple process, ready to accept an external invitation to the parent session on a (resp. to the child session on c) and to behave as expected. P_{bob} is more complex: it sends an invitation on a , and after receiving a request it chooses, as specified in Figure 1, between answering directly on the session channel or contacting the server through a sub-session. In this case, the new session channel k is created and one internal invitation to play role agent in k is sent and accepted by P_{bob} itself, then it proceeds as expected. We describe a reduction sequence for the composition of these three processes:

$$\begin{aligned}
P_{\text{alice}} \mid P_{\text{carol}} \mid P_{\text{bob}} &\rightarrow\rightarrow s?[\text{middle}, \text{client}]\text{answer}(ans_0) \mid P_{\text{carol}} \\
&\mid (\dots) + (\text{new } k \text{ on } s \text{ with } (req_0) \& (c \text{ as server}) \dots) \\
&\rightarrow s?[\text{middle}, \text{client}]\text{answer}(ans_0) \mid P_{\text{carol}} \mid \bar{c}(k) \\
&\mid s \uparrow [\text{middle}, \text{middle} : \text{agent}](k) \mid s \downarrow [\text{middle}, \text{middle} : \text{agent}](z) \dots \\
&\rightarrow s?[\text{middle}, \text{client}]\text{answer}(ans_0) \mid P_{\text{carol}} \\
&\mid \bar{c}(k) \mid k![\text{middle}, \text{server}]\text{request}(req_0) \dots \\
&\rightarrow s?[\text{middle}, \text{client}]\text{answer}(ans_0) \\
&\mid k![\text{agent}, \text{server}]\text{request}(req_0) \dots \\
&\mid k?[\text{agent}, \text{server}]\text{request}(req) \dots \rightarrow\rightarrow\rightarrow \mathbf{0}
\end{aligned}$$

After two communications on a and s , the reduct of P_{bob} reaches the located choice. We suppose it chooses the second branch. Thus, an output on c containing session name k is created. Then the internal self-invitation for k is performed, and, finally, the external invitation of P_{carol} on c . Three reductions can still be played, two on k and one on s .

Validation We describe a static way to ensuring that processes conforms to formal specifications. The *global environment* Γ relates shared channels to the type of the invitation they carry, protocol names to their code and session channels to the global type they implement. $a : T[\mathbf{r}]$ means that a is used to send and receive invitations to play role \mathbf{r} with local type T , $\mathcal{P} : (\tilde{\mathbf{r}}^1, \tilde{\mathbf{y}}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}})$ describes the participants, arguments and code of protocol \mathcal{P} , finally, $s : G$ means that protocol G can be implemented on s . The *session environment* Δ relates pairs of session channels and roles $s[\mathbf{r}]$ to local types. $s[\mathbf{r}] : T$ means that in session s , participant \mathbf{r} still has to perform the actions of T . $s[\mathbf{r}]^\bullet : T$ (resp. $s[\mathbf{r}]^\circ : T$) stands for the *capability* to invite externally (resp. internally) someone to play role \mathbf{r} in s . In the following, we consider only environments which are mappings, and we will write $\Delta(s[\mathbf{r}])$. Additionally, we write $\Delta(s) = \mathbf{0}$ when s does not appear in Δ and $s[\mathbf{r}]^-$ to denote either $s[\mathbf{r}]^\circ$, $s[\mathbf{r}]^\bullet$ or $s[\mathbf{r}]$. We allow "garbage collection" for session environment: $(\Delta, s[\mathbf{r}] : \text{end}) = \Delta$.

$$\begin{aligned}
\Gamma &::= \emptyset \mid \Gamma, a : T[\mathbf{r}] \mid \Gamma, \mathcal{P} : (\tilde{\mathbf{r}}^1, \tilde{\mathbf{y}}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}}) \mid \Gamma, s : G \\
\Delta &::= \emptyset \mid \Delta, s[\mathbf{r}] : T \mid \Delta, s[\mathbf{r}]^\bullet : T \mid \Delta, s[\mathbf{r}]^\circ : T
\end{aligned}$$

A typing judgement $\Gamma \vdash P \triangleright \Delta$ means that under the global environment Γ , the process P is validated by the session environment Δ . We use $\vdash v : S$ to notify that value v has kind S . The validation rules are as follows:

$$\begin{array}{c}
\text{(I, O)} \quad \frac{\Gamma \vdash P \triangleright \Delta, x[\mathbf{r}] : T \quad \Gamma(a) = T[\mathbf{r}]}{\Gamma \vdash a(x).P \triangleright \Delta} \quad \frac{\Gamma \vdash P \triangleright \Delta \quad \Gamma(a) = T[\mathbf{r}]}{\Gamma \vdash \bar{a}(s).P \triangleright \Delta, s[\mathbf{r}]^\bullet : T} \\
\text{(C)} \quad \frac{(\Gamma \vdash P_i \triangleright \Delta, s[\mathbf{r}'] : T_i \quad \vdash y_i : S_i)_{i \in I}}{\Gamma \vdash s?[\mathbf{r}, \mathbf{r}']_{i \in I} \{l_i(y_i).P_i\} \triangleright \Delta, s[\mathbf{r}'] : \mathbf{get}[\mathbf{r}]?_{i \in I} \{l_i(x_i : S_i).T_i\}} \\
\text{(S)} \quad \frac{\Gamma \vdash P \triangleright \Delta, s[\mathbf{r}] : T_j \quad \vdash v : S_j}{\Gamma \vdash s![\mathbf{r}, \mathbf{r}']_{l_j}(v).P \triangleright \Delta, s[\mathbf{r}] : \mathbf{send}[\mathbf{r}]!_{i \in I} \{l_i(x_i : S_i).T_i\}} \\
\text{(P)} \quad \frac{\Gamma \vdash P \triangleright \Delta, s[\mathbf{r}] : T \quad \Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G) \quad G\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}''} = T''}{\Gamma \vdash s \uparrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}''](k).P \triangleright \Delta, s[\mathbf{r}] : \mathbf{req} \mathcal{P}[\mathbf{r}''](\tilde{v}) \mathbf{to} \mathbf{r}'.T, k[\mathbf{r}'']^\circ : T''} \\
\text{(J)} \quad \frac{\Gamma \vdash P \triangleright \Delta, s[\mathbf{r}'].T, x[\mathbf{r}'''] : T'' \quad \Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G) \quad G\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}''} = T''}{\Gamma \vdash s \downarrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}'''](x).P \triangleright \Delta, s[\mathbf{r}'] : \mathbf{ent} \mathcal{P}[\mathbf{r}'''](\tilde{v}) \mathbf{from} \mathbf{r}.T} \\
\text{(New)} \quad \frac{\Gamma \vdash P \triangleright \Delta, s[\mathbf{r}] : T, k[\mathbf{r}_1]^\circ : T'_1, \dots, k^\circ[\mathbf{r}_n] : T'_n, k^\bullet[\mathbf{r}_1^2] : T'_{n+1}, \dots, k^\bullet[\mathbf{r}_m^2] : T'_{n+m} \quad \Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G) \quad \forall i, \Gamma(a_i) = T'_{i+n}[\mathbf{r}_{i+n}] \quad \forall i, G\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}_i^1} = T'_i \quad \forall j, G\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}_j^2} = T'_{j+n} \quad \vdash \tilde{v} : S \quad \Gamma(k) : \mathcal{P}\{\tilde{v}/\tilde{y}\}}{\Gamma \vdash \mathbf{new} \ k \ \mathbf{on} \ s \ \mathbf{with} \ (\tilde{v})\&(\tilde{a} \ \mathbf{as} \ \tilde{\mathbf{r}}^2).P \triangleright \Delta, s[\mathbf{r}] : \mathbf{call} \ \mathcal{P} : G \ \mathbf{with} \ (\tilde{v} \ \mathbf{as} \ \tilde{y} : \tilde{S})\&(\tilde{\mathbf{r}}^2).T} \\
\text{(N, P)} \quad \frac{\Gamma \vdash P_1 \triangleright \Delta_1 \quad \Gamma \vdash P_2 \triangleright \Delta_2}{\Gamma \vdash \mathbf{0} \triangleright \emptyset} \quad \frac{\Gamma \vdash P_1 \triangleright \Delta_1 \quad \Gamma \vdash P_2 \triangleright \Delta_2}{\Gamma \vdash P_1 \mid P_2 \triangleright \Delta_1 \otimes \Delta_2} \\
\text{(S1)} \quad \frac{\Gamma \vdash P_1 \triangleright \Delta, s[\mathbf{r}] : T_1 \quad \Gamma \vdash P_2 \triangleright \Delta, s[\mathbf{r}] : T_2}{\Gamma \vdash P_1 + P_2 \triangleright \Delta, s[\mathbf{r}] : T_1 \oplus T_2} \\
\text{(S2, R)} \quad \frac{\Gamma \vdash P \triangleright \Delta, s[\mathbf{r}] : T_i \quad i \in \{1, 2\}}{\Gamma \vdash P \triangleright \Delta, s[\mathbf{r}] : T_1 \oplus T_2} \quad \frac{\Gamma, a : T[\mathbf{r}] \vdash P \triangleright \Delta}{\Gamma \vdash (\nu a) P \triangleright \Delta}
\end{array}$$

Rule **(New)** is the crux of this type system, as it ensures subsessions are called in a sound way. To type the process $\mathbf{new} \ k \ \mathbf{on} \ s \ \mathbf{with} \ (\tilde{v})\&(\tilde{a} \ \mathbf{as} \ \tilde{\mathbf{r}}^2).P$, the session channel k should be associated with a protocol $G\{\tilde{v}/\tilde{y}\}$ matching the one present in the local type of \mathbf{r} in the parent session s : $\mathbf{call} \ \mathcal{P} : G \ \mathbf{with} \ (\tilde{v} \ \mathbf{as} \ \tilde{y} : \tilde{S})\&(\tilde{\mathbf{r}}^2).T$ and global environment Γ should map \mathcal{P} to $(\tilde{\mathbf{r}}^1, \tilde{y}; \tilde{\mathbf{r}}^2; G)$. The endpoint projections $(T_p)_{1 \leq p \leq n+m}$ of \mathcal{P} are divided into two sets, the ones that correspond to roles $(\mathbf{r}_i^1)_{1 \leq i \leq n}$ internally invited, and the ones that correspond to roles $(\mathbf{r}_j^2)_{1 \leq j \leq m}$ externally invited through \tilde{a} . Capabilities $(k[\mathbf{r}_i]^\circ : T_i)_i, (k[\mathbf{r}_j]^\bullet : T_j)_j$ for both types of invitations are given to the continuation process P . Rule **(P)** types a process whose role \mathbf{r} on session s consists in sending an internal invitation to play role \mathbf{r}'' in session k . The process is required to hold the capability for $k[\mathbf{r}'']$, we ensure it corresponds to the type of the invitation. Its counterpart **(J)** ensures that role \mathbf{r} of session s after receiving an invitation for $k[\mathbf{r}'']$, gets the corresponding local type T'' in its Δ . Rules **(I)** and **(O)** handle external invitations. As in the internal case, we ensure that the sending process has the corresponding capability. Rules **(C)** and **(S)** address branching communications on session channels. In both rules we ensure that the values communicated x_i, v have the same value-type as the identifiers y_i in the type. Summations are handled by two rules **(S1)** and **(S2)**. If the local type specifies a choice between two branches, the process can either implement this choice with the $+$ constructor, or implement only one branch of the choice. This illustrates the fact that the decision can be made at implementation time (for instance a middleware implementing *CMS*

which always contacts the server) or at run-time (a middleware which can proceed both ways according to the request). Rule **(Pa)** requires a small explanation, as it allows one to split local types into two branches. We define the \otimes operator with $\Delta_1 \otimes \emptyset = \Delta_1$, $\Delta_1 \otimes (\Delta_2, s[\mathbf{r}]^- : T) = (\Delta_1, s[\mathbf{r}]^- : T) \otimes \Delta_2$ if $\Delta_1(s[\mathbf{r}]) = 0$ and $(\Delta_1, s[\mathbf{r}]^- : T_1) \otimes (\Delta_2, s[\mathbf{r}]^- : T_2) = (\Delta_1, s[\mathbf{r}]^- : T_1 \parallel T_2) \otimes \Delta_2$. Thus, when splitting the session environment in a parallel constructor, we allow the splitting of a single local type composed of two parallel subtypes. Finally, rule **(N)** specifies that the session environment should be empty to type $\mathbf{0}$. This ensures that the processes eventually complete the local types of their specification.

Following the typing rules, one can type the processes introduced in Section 3 as follows: $\Gamma \vdash P_{\text{alice}} \triangleright \emptyset$, $\Gamma \vdash P_{\text{bob}} \triangleright s[\text{middle}] : T_{\text{middle}}, s[\text{client}]^\bullet : T_{\text{client}}$, $\Gamma \vdash P_{\text{carol}} \triangleright \emptyset$ with $\Gamma = a : T_{\text{client}}, c : T_{\text{server}}, \text{Contact} : (\text{agent}, \text{req}; \text{server}; G_{\text{Contact}}), s : G_{\text{Contact}}, k : \text{Contact}, G_{\text{Contact}} \Downarrow_{\text{server}}^\emptyset = T_{\text{server}}$, and $T_{\text{client}} = G_{\text{CMS}} \Downarrow_{\text{client}}^\emptyset, T_{\text{middle}} = G_{\text{CMS}} \Downarrow_{\text{middle}}^\emptyset$ as defined in Section 2.

Session environments for P_{alice} and P_{carol} are empty: processes are not bound to do anything as long as they did not receive an invitation. Session environment for P_{bob} contains both the local type for the role middle played by the process and the capability to send an external invitation for client in the same session. The capability to send an external invitation to server is not created yet.

4 Properties

In this section we justify our theory with two main propositions, subject reduction and progress. First, we define a satisfaction relation relating dynamically processes and specifications. We introduce Labelled Transition Systems for both the processes and the specification. Labels are defined by $\ell ::= \tau \mid \bar{a}\langle v \rangle \mid a\langle v \rangle \mid s?[\mathbf{r}, \mathbf{r}']l\langle k \rangle \mid s![\mathbf{r}, \mathbf{r}']l\langle k \rangle \mid s \downarrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}'']\langle k \rangle \mid s \uparrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}'']\langle k \rangle$. The subject of a label $\mathbf{sbj}(\ell)$ is defined intuitively for all labels, knowing that $\mathbf{sbj}(\tau) = 0$. Labels $\bar{a}\langle v \rangle, a\langle v \rangle, s![\mathbf{r}, \mathbf{r}']l\langle k \rangle, s \uparrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}'']\langle k \rangle$ and τ (resp. $s?[\mathbf{r}, \mathbf{r}']l\langle k \rangle$ and $s \downarrow [\mathbf{r}, \mathbf{r}' : \mathbf{r}'']\langle k \rangle$) are denoted as *output labels* (resp. *input labels*). In the satisfaction relation defined below, output labels are the ones played by the process, to which the specification must answer (thus τ and $a\langle v \rangle$ are considered outputs), and the input labels are the ones the specification plays, to which the process must answer. Transitions for processes $P \xrightarrow{\ell} P'$ follow the reduction semantics. The most relevant transitions for specifications, defined w.r.t. a global environment Γ , are presented in Figure 2.

Definition 2 (Satisfaction) *We say that \mathcal{R}_Γ is a satisfaction relation between process P and specification Δ , if:*

whenever $\Delta \xrightarrow[\Gamma]{\ell} \Delta'$ with an input label ℓ , then $P \xrightarrow{\ell} P'$ and $P' \mathcal{R}_\Gamma \Delta'$,

whenever $P \xrightarrow{\ell} P'$ with an output label ℓ , then $\Delta \xrightarrow[\Gamma]{\ell} \Delta'$ and $P' \mathcal{R}_\Gamma \Delta'$.

The largest relation \mathcal{R}_Γ is called satisfaction w.r.t. Γ denoted $\mathbf{sat}(P, \Delta)_\Gamma$. In this case, we say that P satisfies Δ w.r.t. Γ (we omit this last part when Γ is clear from context)

$$\begin{array}{c}
\Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y} : \tilde{S}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}}) \quad \forall i, G_{\mathcal{P}}\{\tilde{v}/\tilde{y}\} \Downarrow_{\tilde{\mathbf{r}}^1_i} = T'_i \\
\forall j, G_{\mathcal{P}}\{\tilde{v}/\tilde{y}\} \Downarrow_{\tilde{\mathbf{r}}^2_j} = T''_j \quad \Gamma(k) = G\{\tilde{v}/\tilde{y}\} \\
\text{(Ssub)} \frac{}{s[\mathbf{x}] : \text{call } \mathcal{P} : G \text{ with } (\tilde{v} \text{ as } \tilde{y} : \tilde{S}) \& (\tilde{\mathbf{r}}^2). T \xrightarrow{\tau} s[\mathbf{x}] : T, (k[\mathbf{r}^1_i] : T'_i)_i, (k[\mathbf{r}^2_j] : T''_j)_j} \\
\text{(Sout)} \frac{\Gamma(a) = T[\mathbf{x}]}{k[\mathbf{x}]^\bullet : T \xrightarrow{\Gamma} \emptyset} \quad \text{(Sin)} \frac{\Gamma(a) = T[\mathbf{x}]}{\emptyset \xrightarrow{\Gamma} k[\mathbf{x}] : T} \quad \text{(ScomC)} \frac{}{k[\mathbf{r}''^\bullet] : T'' \xrightarrow{\Gamma} k[\mathbf{r}'''] : T''} \\
\text{(Sjoin)} \frac{\Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y} : \tilde{S}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}}) \quad G_{\mathcal{P}}\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}''} = T''}{s[\mathbf{r}'] : \text{ent } \mathcal{P}[\mathbf{r}''](\tilde{v}) \text{ from } \mathbf{r}. T \xrightarrow{\Gamma} s[\mathbf{r}'] : T, k : [\mathbf{r}'''] : T''} \\
\text{(Sparti)} \frac{\Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y} : \tilde{S}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}}) \quad G_{\mathcal{P}}\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}''} = T''}{s[\mathbf{x}] : \text{req } \mathcal{P}[\mathbf{r}''](\tilde{v}) \text{ to } \mathbf{r}'. T, k[\mathbf{r}''^\circ] : T'' \xrightarrow{\Gamma} s[\mathbf{x}] : T} \\
\text{(Sinvit)} \frac{\Gamma(\mathcal{P}) = (\tilde{\mathbf{r}}^1, \tilde{y} : \tilde{S}; \tilde{\mathbf{r}}^2; G_{\mathcal{P}}) \quad G_{\mathcal{P}}\{\tilde{v}/\tilde{y}\} \Downarrow_{\mathbf{r}''} = T''}{s[\mathbf{x}] : \text{ent } \mathcal{P}[\mathbf{r}''](\tilde{v}) \text{ from } \mathbf{r}. T, s[\mathbf{r}'] : \text{req } \mathcal{P}[\mathbf{r}''](\tilde{v}) \text{ to } \mathbf{r}. T', k[\mathbf{r}''^\circ] : T'' \xrightarrow{\Gamma} s[\mathbf{x}] : T, s[\mathbf{r}'] : T', k[\mathbf{r}'''] : T''}
\end{array}$$

Fig. 2. Transitions for specifications (excerpt)

We justify the soundness of our framework by relating the static validation to the dynamic satisfaction, through *correspondence*. If a process is validated by a specification Δ , it is able to behave as described in Δ . From this property, we derive subject reduction, which ensures that validation is preserved by reduction. A session environment is *coherent* if it is composed of projections of well-formed global types. A coherent session environment is *simple* if it consists of a single session. A process is *unblocked* if it does not contain hidden channels and if its session channel is never under a prefix whose subject is a shared channel, except when the latter binds the former. If an unblocked process is validated by a simple coherent session environment, interactions at session channels can proceed. If, further, the original global type is non-recursive, the process can eventually complete all interactions at its session-environment.

Proposition 3 (Soundness of the type system)

(*Correspondence*) If $\Gamma \vdash P \triangleright \Delta$ then $\text{sat}(P, \Delta)_\Gamma$.

(*Subject Reduction*) If $\Gamma \vdash P \triangleright \Delta$ and $P \rightarrow P'$ then there exists Δ' s.t. $\Gamma \vdash P' \triangleright \Delta'$.

(*Progress*) If P is unblocked and $\Gamma \vdash P \triangleright \Delta$ such that Δ is simple, then there exists P' s.t. $P \rightarrow^+ P'$, $\Gamma \vdash P' \triangleright \Delta'$ and Δ' is coherent.

(*Coherence*) If P is unblocked and $\Gamma \vdash P \triangleright \Delta$ such that Δ is simple, and moreover Δ does not contain recursions, then there exists P' s.t. $P \rightarrow^* P'$ and $\Gamma \vdash P' \triangleright \emptyset$.

5 Returning a result

We introduce the notion of *result* of a session as an object (which can be a value or even a protocol), sent back to the initiator of the session. *Protocols with results* allow us to describe complex governance properties, such as ensuring that a privately negotiated price corresponds to the one proposed publicly in the parent protocol. Suppose we want to ensure that, in *CMS*, the answer *ans* given in the subprotocol *Contact* by server is the same as *ans₀* sent by middle to client. Information can be transmitted from a parent session to a subsession, but the converse is not possible. Continuation-Passing-Style is a possible solution: we convert the end of the *CMS* protocol into a continuation *K*, send it as argument when calling *Contact* and call *K* inside *Contact* with *ans*. However, this may not lead to a clean descriptive framework. Thus we choose to use a dedicated mechanism. The syntax of global types with results adds **r returns**(*res* : *S*) and (*res* : *S*) ← **r calls** $\mathcal{P}(\tilde{x}, \tilde{v})$ (replacing **end** and **r calls** $\mathcal{P}(\tilde{v})$). The former constructor ends the session by specifying that the protocol returns the value identified in the session by *res* and that **r** is responsible for doing it, the latter specifies that we call a subprotocol which eventually produces a result *res*. Kinds ensure that the returned result has the type expected by the initiator.

We present corresponding modifications to *CMS*. Inside the *Contact* protocol, we ask agent to send the result *ans* back to the parent protocol. In the latter, the result *ans₀* is expected when calling *Contact*, thus we ensure that the answer sent by the server in the subprotocol is the same as the one sent to client in the parent one. Local types use similar constructors and implementation of result is done through cross-session communications.

```
let Contact = (agent, req : Req){
... agent returns(ans : Req)
in
... ans0 : Req ← middle calls Contact(middle, req0) ...
```

In the framework presented above, subsessions are executed in parallel with the parent session. The result mechanism allows one to include synchronisation between the two sessions:

```
... Alice waits for contract calling Nego(Alice, Carol).
Alice → Bob : Data(contract) ...
```

Here participant Alice starts a negotiation subsession with Carol. When the negotiation is over, she sends the result of the subsession to Bob, participant of the parent session not invited in the subsession. This has two advantages, first Bob can know the result of a subsession without going through the internal invitation procedure, and it prevents both Alice and Bob to perform actions in the parent session as long as the subsession is not over.

6 Conclusion and Future Works

To our knowledge, there does not exist other works addressing the notion of nested session types, or protocol calls inside session types. The closest contribu-

tion is [9], which introduces parametrisation of protocols through *dynamic session types*. Parametrisation allows one single two-party protocol to be applied to each pair of agents in a large network. Our framework contains more than simple parametrisation, it presents nesting and introduces kinds and higher-order programming. Another related work is [7]: the authors describe a global language for *choreographies*, implementing global types, protocols interleaved in the same choreography can be merged together into a single global type, removing costly invitations. The authors actually proceed in a direction different from ours, by trying to unify every protocol into a single superprotocol. Their approach focuses on implementation, while ours focus on types. We believe session type theory benefits independently from both methods. Our contribution makes use of the same formal framework as [3, 13, 2]. Each of these contributions adds expressiveness, in different directions (logical assertions, ghost states, monitoring), to a large common theory for validation of distributed programs with session types. The whole theory (including this work) is put in practise by the development of the Scribble language [16] and the collaboration with the OOI project [15].

We are currently investigating how the result mechanism can be improved (in the context of [16]). Currently, the result is sent to the initiator. Broadcasting the result to every member of the subsession might also be a desirable feature. Moreover, our results are restricted to value-types, but some use cases of [15] specify that a negotiation subprotocol produces a contract that is used in the parent protocol to control interactions. Although it leads to technical challenges, we believe our framework can eventually accommodate such behaviours by using dependent types, introducing abstract logical predicate decided at run-time inside global types. Exceptions handling in a distributed asynchronous framework, remains a challenging task, even if some progress have been made in [6] and [5]. Yet exceptions are absolutely necessary when specifying real-world protocols. We believe that nested protocols give a simple way to handle exceptions, by making explicit blocks of computation.

Acknowledgements We thank the CONCUR reviewers for their comments, our colleagues in Mobility Reading Group for discussions, and the OOI project and Matthew Arrott for their feedback. This work is supported by Ocean Observatories Initiative [15] and EPSRC grants EP/F002114/1 and EP/G015481/1.

References

1. Lorenzo Bettini, Mario Coppo, Loris D’Antoni, Marco De Luca, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. Global progress in dynamically interleaved multiparty sessions. In *CONCUR*, pages 418–433, 2008.
2. Laura Bocchi, Pierre-Malo Denéilou, Romain Demangeon, Kohei Honda, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. Dynamic and static safety validation in distributed programs through multiparty sessions. (submitted), 2012.
3. Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. A theory of design-by-contract for distributed multiparty interactions. In *CONCUR*, volume 6269 of *LNCS*, pages 162–176, 2010.

4. Sara Capecchi, Ilaria Castellani, and Mariangiola Dezani-Ciancaglini. Information flow safety in multiparty sessions. In *EXPRESS*, volume 64 of *EPTCS*, pages 16–30, 2011.
5. Sara Capecchi, Elena Giachino, and Nobuko Yoshida. Global escape in multiparty sessions. In Kamal Lodaya and Meena Mahajan, editors, *FSTTCS*, volume 8 of *LIPICs*, pages 338–351. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.
6. Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured interactional exceptions in session types. In *CONCUR*, volume 5201 of *LNCS*, pages 402–417. Springer, 2008.
7. Marco Carbone and Fabrizio Montesi. Merging multiparty protocols in multiparty choreographies. (unpublished, presented at PLACES), 2012.
8. Mario Coppo and Mariangiola Dezani-Ciancaglini. Structured communications with concurrent constraints. In *TGC*, pages 104–125, 2008.
9. Pierre-Malo Denielou and Nobuko Yoshida. Dynamic multirole session types. In *POPL*, pages 435–446, 2011.
10. Mariangiola Dezani-Ciancaglini. A reputation system for multirole sessions. Invited talk at TGC, September 2011.
11. Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP'98*, volume 1381 of *LNCS*, pages 22–138. Springer, 1998.
12. Ivan Lanese, Claudio Guidi, Fabrizio Montesi, and Gianluigi Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In Antonio Cerone and Stefan Gruner, editors, *SEFM*, pages 323–332. IEEE Computer Society, 2008.
13. Romain Demangeon Laura Bocchi and Nobuko Yoshida. A multiparty multi-session logic. (submitted), 2012.
14. Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, Berlin, 1980.
15. Ocean Observatories Initiative (OOI). <http://www.oceanleadership.org/programs-and-partnerships/ocean-observing/ooi/>.
16. Scribble Project homepage. www.scribble.org.