

Automatic Generation of Self-Monitoring MASs from Multiparty Global Session Types in Jason

Davide Ancona¹, Sophia Drossopoulou², and Viviana Mascardi¹

¹ DISI, University of Genova, Italy
{davide.ancona,viviana.mascardi}@unige.it

² Imperial College, London, UK
scd@doc.ic.ac.uk

Abstract. Global session types are behavioral types designed for specifying in a compact way multiparty interactions between distributed components, and verifying their correctness. We take advantage of the fact that global session types can be naturally represented as cyclic Prolog terms - which are directly supported by the Jason implementation of AgentSpeak - to allow simple automatic generation of self-monitoring MASs: given a global session type specifying an interaction protocol, and the implementation of a MAS where agents are expected to be compliant with it, we define a procedure for automatically deriving a self-monitoring MAS. Such a generated MAS ensures that agents conform to the protocol at run-time, by adding a *monitor* agent that checks that the ongoing conversation is correct w.r.t. the global session type. The feasibility of the approach has been experimented in Jason for a non-trivial example involving recursive global session types with alternative choice and fork type constructors. Although the main aim of this work is the development of a unit testing framework for MASs, the proposed approach can be also extended to implement a framework supporting self-recovering MASs.

1 Introduction

A protocol represents an agreement on how participating systems interact with each other. Without a protocol, it is hard to do a meaningful interaction: participants simply cannot communicate effectively.

The development and validation of programs against protocol descriptions could proceed as follows:

- *A programmer specifies a set of protocols to be used in her application.*

...

- *At the execution time, a local monitor can validate messages with respect to given protocols, optionally blocking invalid messages from being delivered.*

This paper starts with a few sentences drawn from the manifesto of Scribble, a language to describe application-level protocols among communicating systems

initially designed by Kohei Honda and Gary Brown¹. The team working on Scribble involves both scientists active in the agent community and scientists active in the session types one. Their work inspired the proposal presented in this paper where multiparty global session types are used on top of the Jason agent oriented programming language for runtime verification of the conformance of a MAS implementation to a given protocol. This allows us to experiment our approach on realistic scenarios where messages may have a complex structure, and their content may change from one interaction to another.

Following Scribble’s manifesto, we ensure runtime conformance thanks to a Jason monitor agent that can be automatically generated from the global session type, represented as a Prolog cyclic term. Besides the global session type, the developer must specify the type of the actual messages that are expected to be exchanged during a conversation.

In order to verify that a MAS implementation is compliant with a given protocol, the Jason code of the agents that participate in the protocol is extended seamlessly and automatically. An even more transparent approach would be possible by overriding the underlying agent architecture methods of Jason responsible for sending and receiving messages, which could intercept all messages sent by the monitored agents, and send them to the monitor which could manage them in the most suitable way. In this approach message “sniffing” would have to occur at the Java (API) level, gaining in transparency but perhaps losing in flexibility.

In this paper we show the feasibility of our approach by testing a MAS against a non-trivial protocol involving recursive global session types with alternative choice and fork type constructors.

The paper is organized in the following way: Section 2 provides a gentle introduction to the notion of global session type adopted in this work; Section 3 discusses our implementation of the protocol testing mechanism and presents the results of our experiments; Section 4 discusses the related literature and outlines future directions of our work.

2 A gentle introduction to global session types for agents

In this section we informally introduce global session types (global types for short) and show how they can be smoothly integrated in MASs to specify multiparty communication protocols between agents. To this aim, we present a typical protocol that can be found in literature as our main running example used throughout the paper.

Our example protocol involves three different agents playing the roles of a seller s , a broker b , and a client c , respectively. Such a protocol is described by the FIPA AUML interaction diagram [14] depicted in Figure 1: initially, s communicates to b the intention to sell a certain item to c ; then the protocol enters a negotiation loop of an arbitrary number n (with $n \geq 0$) of iterations,

¹ <http://www.jboss.org/scribble/>

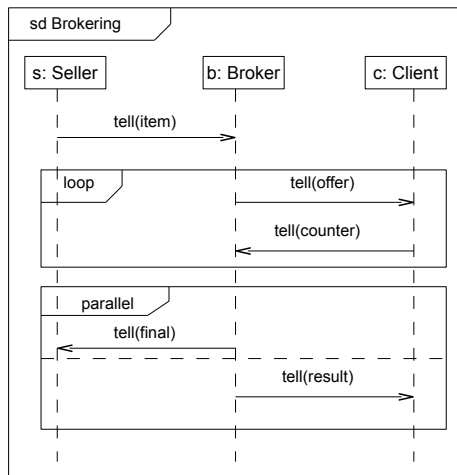


Fig. 1. The Brokering interaction protocol in FIPA AUML.

where **b** sends an offer to **c** and **c** replies with a corresponding counter-offer. After such a loop, **b** concludes the communication by sending in an arbitrary order the message of type **result** to **c**, and of type **final** to **s**.

Even though the AUML diagram of Figure 1 is very intuitive and easy to understand, a more compact and formal specification of the protocol is required to perform verification or testing of a MAS, in order to provide guarantees that the protocol is implemented correctly. Global session types [5, 11] have been introduced and studied exactly for this purposes, even though in the more theoretical context of calculi of communicating processes. A global type describes succinctly all sequences of sending actions that may occur during a correct implementation of a protocol.

Depending on the employed type constructors, a global type can be more or less expressive. Throughout this paper we will use a fixed notion of global type, but our proposed approach can be easily adapted for other kinds of global types. The notion of global type we adopt is a slightly less expressive version of that proposed by Deniélou and Yoshida [7] (which, however, allows us to specify the protocol depicted in Figure 1), defined on top of the following type constructors:

- *Sending Actions*: a sending action occurs between two agents, and specifies the sender and the receiver of the message (in our case, the names of the agents, or, more abstractly, the role they play in the communication), and the type of the performative and of the content of the sent message; for

instance, `msg(s, b, tell, item)` specifies that agent `s` (the seller) sends the `tell` performative to agent `b` (the broker) with content of type `item`.

- *Empty Type*: the constant `end` represents the empty interaction where no sending actions occur.
- *Sequencing*: sequencing is a binary constructor allowing a global type t to be prefixed by a sending action a ; that is, all valid sequences of sending actions denoted by $\text{seq}(a, t)$ are obtained by prefixing with a all those sequences denoted by t . For instance,

```
seq(msg(alice, bob, tell, ping),
    seq(msg(bob, alice, tell, pong), end))
```

specifies the simple interaction where first `alice` sends `tell(ping)` to `bob`, then `bob` replies to `alice` with `tell(pong)`, and finally the interaction stops.

- *Choice*: the choice constructor has variable arity² n (with $n \geq 0$) and expresses an alternative between n possible choices. Because its arity is variable we use a list to represent its operands. For instance,

```
choice([
    seq(msg(c, b, tell, counter), end),
    seq(msg(b, s, tell, final), end),
    seq(msg(b, c, tell, result), end)
])
```

specifies an interaction where either `c` sends `tell(counter)` to `b`, or `b` sends `tell(final)` to `s`, or `b` sends `tell(result)` to `c`.

- *Fork*: the fork binary³ constructor specifies two interactions that can be interleaved. For instance,

```
fork(
    seq(msg(b, s, tell, final), end),
    seq(msg(b, c, tell, result), end)
)
```

specifies the interaction where first `b` sends `tell(final)` to `s`, and then `b` sends `tell(result)` to `c`, or the other way round.

Recursive types: the example types shown so far do not specify any interaction loop, as occurs in the protocol of Figure 1. To specify loops we need to consider recursive global types; for instance, the protocol consisting of infinite sending actions where first `alice` sends `tell(ping)` to `bob`, and then `bob` replies `tell(pong)` to `alice`, can be represented by the recursive type T s.t.

```
T = seq(msg(alice, bob, tell, ping),
        seq(msg(bob, alice, tell, pong), T))
```

If we interpret the equation above syntactically (that is, as a unification problem), then the unique solution is an infinite term (or, more abstractly, an infinite

² Arity 0 and 1 are not necessary, but make the definition of predicate `next` simpler.

³ For simplicity, the operator has a fixed arity, but it could be generalized to the case of n arguments (with $n \geq 2$) as happens for the `choice` constructor.

tree) which is *regular*, that is, whose set of subterms is finite. In practice, the unification problem above is solvable in most modern implementations of Prolog, where cyclic terms are supported; this happens also for the Jason implementation, where Prolog-like rules can be used to derive beliefs that hold in the current belief base⁴. As another example, let us consider the type `T2` s.t.

```
T2 = seq(msg(alice,bob,tell,ping),
         seq(msg(bob,alice,tell,pong),choice([T2,end])))
```

Such a type contains the infinite interaction denoted by `T` above, but also all finite sequences of length $2n$ (with $n \geq 1$) of alternating sending actions `msg(alice,bob,tell,ping)` and `msg(bob,alice,tell,pong)`.

We are now ready to specify the **Brokering** protocol with a global type `BP`, where for sake of clarity we use the auxiliary types `OffOrFork`, `Off`, and `Fork`:

```
BP      = seq(msg(s,b,tell,item),OffOrFork),
OffOrFork = choice([Off,Fork])
Off     = seq(msg(b,c,tell,offer),
              seq(msg(c,b,tell,counter),OffOrFork))
Fork    = fork(seq(msg(b,s,tell,final),end),
              seq(msg(b,c,tell,result),end))
```

Note that for the definition of global types we consider in this paper, the `fork` constructor does not really extend the expressiveness of types: any type using `fork` can be transformed into an equivalent one without `fork`. However, such a transformation may lead to an exponential growth of the type [2].

Formal definitions

Figure 2 defines the abstract syntax of the global session types that will be used in the rest of the paper. As already explained in the previous section, global

$$\begin{aligned} \text{GT} ::= & \text{choice}([\text{GT}_1, \dots, \text{GT}_n]) \ (n \geq 0) \mid \\ & \text{seq}(\text{SA}, \text{GT}) \mid \\ & \text{fork}(\text{GT}_1, \text{GT}_1) \mid \\ & \text{end} \\ \text{SA} ::= & \text{msg}(\text{AId}_1, \text{AId}_2, \text{PE}, \text{CT}) \end{aligned}$$

Fig. 2. Syntax of Global Types.

types are defined coinductively: `GT` is the greatest set of regular terms defined by the productions of Figure 2.

The meta-variables `AId`, `PE` and `CT` range over agent identifiers, performatives, and content types, respectively. Content types are constants specifying the types of the contents of messages.

⁴ Persistency of cyclic terms is supported by the very last version of Jason; since testing of this feature is still ongoing, it has not been publicly released yet.

The syntactic definition given so far still contains global types that are not considered useful, and, therefore, are rejected for simplicity. Consider for instance the following type `NC`:

```
NC = choice([NC,NC])
```

Such a type is called *non contractive* (or *non guarded*), since it contains an infinite path with no `seq` type constructors. These kinds of types pose termination problems during dynamic global typechecking. Therefore, in the sequel we will consider only *contractive* global types (and we will drop the term “contractive” for brevity), that is, global types that do not have paths containing only the `choice` and `fork` type constructors. Such a restriction does not limit the expressive power of types, since it can be shown that for every non contractive global type, there exists a contractive one which is equivalent, in the sense that it represents the same set of sending action sequences. For instance, the type `NC` as defined above corresponds to the empty type `end`.

Interpretation of global types. We have already provided an intuition of the meaning of global types. We now define their interpretation, expressed in terms of a `next` predicate, specifying the possible transitions of a global type. Intuitively, a global type represents a state from which several transition steps to other states (that is, other global types) are possible, with a resulting sending action. Consider for instance the type `F` defined by

```
fork(seq(msg(b,s,tell,final),end),
      seq(msg(b,c,tell,result),end))
```

Then there are two possible transition steps: one yields the sending action `msg(b,s,tell,final)` and moves to the state corresponding to the type

```
fork(end,
      seq(msg(b,c,tell,result),end))
```

while the other yields the sending action `msg(b,c,tell,result)` and moves to the state corresponding to the type

```
fork(seq(msg(b,s,tell,final),end),
      end)
```

Predicate `next` is defined below, with the following meaning: if `next(GT1,SA,GT2)` succeeds, then there is a one step transition from the state represented by the global type `GT1` to the state represented by the global type `GT2`, yielding the sending action `SA`. The predicate is intended to be used with the mode indicators `next(+,+,-)`, that is, the first two arguments are input, whereas the last is an output argument.

```
1 next(seq(msg(S, R, P, CT),GT),msg(S, R, P, C),GT) :-
   has_type(C, CT).
2 next(choice([GT1|_]),SA,GT2) :- next(GT1,SA,GT2).
3 next(choice([_|L]),SA,GT) :- next(choice(L),SA,GT).
4 next(fork(GT1,GT2),SA,fork(GT3,GT2)) :- next(GT1,SA,GT3).
5 next(fork(GT1,GT2),SA,fork(GT1,GT3)) :- next(GT2,SA,GT3).
```

We provide an explanation for each clause:

1. For a sequence $\text{seq}(\text{msg}(\text{S}, \text{R}, \text{P}, \text{CT}), \text{GT})$ the only allowed transition step leads to state GT , and yields a sending action $\text{msg}(\text{S}, \text{R}, \text{P}, \text{C})$ where C is required to have type CT ; we assume that all used content types are defined by the predicate has_type , whose definition is part of the specification of the protocol, together with the initial global type.
2. The first clause for choice states that there exists a transition step from $\text{choice}([\text{GT1}|_])$ to GT2 yielding the sending action SA , whenever there exists a transition step from GT1 to GT2 yielding the sending action SA .
3. The second clause for choice states that there exists a transition step from $\text{choice}([_|\text{L}])$ to GT yielding the sending action SA , whenever there exists a transition step from $\text{choice}(\text{L})$ (that is, the initial type where the first choice has been removed) to GT yielding the sending action SA .
Note that both clauses for choice fail for the empty list, as expected (since no choice can be made).
4. The first clause for fork states that there exists a transition from $\text{fork}(\text{GT1}, \text{GT2})$ to $\text{fork}(\text{GT3}, \text{GT2})$ yielding the sending action SA , whenever there exists a transition step from GT1 to GT3 yielding the sending action SA .
5. The second clause for fork is symmetric to the first one.

We conclude this section by a claim stating that contractive types ensure termination of the resolution of next .

Proposition 1. *Let us assume that $\text{has_type}(c, ct)$ always terminates for any ground atoms c and ct . Then, $\text{next}(gt, sa, X)$ always terminates, for any ground terms gt and sa , and logical variable X , if gt is a contractive global type.*

Proof. By contradiction, it is straightforward to show that if $\text{next}(gt, sa, X)$ does not terminate, then gt must contain a (necessarily infinite) path with only choice and fork constructors, hence, gt is not contractive.

3 A Jason Implementation of a Monitor for Checking Global Session Types

As already explained in the Introduction, the main motivation of our work is a better support for testing the conformance of a MAS to a given protocol, even though we envisage other interesting future application scenarios (see Section 4). From this point of view our approach can be considered as a first step towards the development of a unit testing framework for MASs where testing, types, and – more generally – formal verification can be reconciled in a synergistic way.

In more detail, given a Jason implementation of a MAS⁵, our approach allows automatic generation⁶ of an extended MAS from it, that can be run on a set of tests to detect possible deviations of the behavior of a system from a given

⁵ We assume that the reader is familiar with the AgentSpeak language [17].

⁶ Its implementation has not been completed yet.

protocol. To achieve this the developer is required to provide (besides the original MAS, of course) the following additional definitions:

- The Prolog clauses for predicate `next` defining the behavior of the used global types (as shown in Section 2); such clauses depend on the notion of global type needed for specifying the protocol; depending on the complexity of the protocol, one may need to adopt more or less expressive notions of global types, containing different kinds of type constructors, and for each of them the corresponding behavior has to be defined in terms of the `next` predicate. However, we expect the need for changing the definition of `next` to be a rare case; the notion of global type we present here captures a large class of frequently used protocols, and it is always possible to extend the testing unit framework with a collection of predefined notions of global types among which the developer can choose the most suitable one.
- The global type specifying the protocol to be tested; this can be easily defined in terms of a set of unification equations.
- The clauses for the `has_type` predicate (already mentioned in Section 2), defining the types used for checking the content of the messages; also in this case, a set of predefined primitive types could be directly supported by the framework, leaving to the developer the definition of the user-defined types.

The main idea of our approach relies on the definition of a centralized *monitor* agent that verifies that a conversation among any number of participants is compliant with a given global type, and warns the developer if the MAS does not progress. Furthermore, the code of the agents of the original MAS requires minimal changes that, however, can be performed in an automatic way.

In the sequel, we describe the code of the monitor agent, and the changes applied to all other agents (that is, the participants of the implemented protocol).

3.1 Monitor

We illustrate the code for the monitor by using our running brokering example. The monitor can be automatically generated from the global type specification in a trivial way. The global type provided by the developer is simply a conjunction *UnifEq* of unification equations of the form $X = GT$, where X is a logical variable, and GT is a term (possibly containing logical variables) denoting a global type. The use of more logical variables is allowed for defining auxiliary types that make the definition of the main type more readable. Then from *UnifEq* the following Prolog rule is generated:

```
initial_state(X) :- UnifEq.
```

where X is the logical variable contained in *UnifEq* corresponding to the main global type. The definition of the type of each message content must be provided as well. In fact, the protocol specification defines also the expected types (such as `item`, `offer`, `counter`, `final` and `result`) for the correct content of all possible messages. For example, the developer may decide that the type `offer` defines all

terms of shape `offer(Item, Offer)`, where `Item` is a string and `Offer` is an integer; similarly, the type `item` corresponds to all terms of shape `item(Client, Item)` where both `Client` and `Item` are strings.

Consequently, the developer has to provide the following Prolog rules that formalize the descriptions given above:

```
has_type(offer(Item, Offer), offer) :-
    string(Item) & int(Offer).
has_type(item(Client, Item), item) :-
    string(Client) & string(Item).
```

The monitor keeps track of the runtime evolution of the protocol by saving its current state (corresponding to a global type), and checking that each message that a participant would like to send, is allowed by the current state. If so, the monitor allows the participant to send the message by explicitly sending an acknowledgment to it. We explain how participants inform the monitor of their intention to send a message in Section 3.2.

The correctness of a sending action is directly checked by the `next` predicate, that also specifies the next state in case the transition is correct. In other words, verifying the correctness of the message sent by `S` to `R` with performative `P` and content `C` amounts to checking if it is possible to reach a `NewState` from the `CurrentState`, yielding a sending action `msg(S, R, P, C)` (`type_check` predicate).

```
/* Monitor's initial beliefs and rules */

// user-defined predicates
initial_state(Glob) :-
    Merge = choice([Off,Fork]) &
    Off = seq(msg(b, c, tell, offer),
              seq(msg(c, b, tell, counter), Merge)) &
    Fork = fork(seq(msg(b, s, tell, final),end),
                seq(msg(b, c, tell, result),end)) &
    Glob = seq(msg(s, b, tell, item),Merge).

has_type(offer(Item, Offer), offer) :-
    string(Item) & int(Offer).
has_type(counter(Item, Offer), counter) :-
    string(Item) & int(Offer).
has_type(final(Res, Client, Item, Offer), final) :-
    string(Res) & string(Client) & string(Item) & int(Offer).
has_type(result(Res, Item, Offer), result) :-
    string(Res) & string(Item) & int(Offer).
has_type(item(Client, Item), item) :-
    string(Client) & string(Item).
// end of user-defined predicates

timeout(4000).

type_check(msg(S, R, P, C), NewState) :-
    current_state(CurrentState) &
    next(CurrentState, msg(S, R, P, C), NewState).

// Rules defining the next predicate follow
.....
```

The monitor prints every information relevant for testing on the console with the `.print` internal action. The `.send(R, P, C)` internal action implements the

asynchronous delivery of a message with performative **P** and content **C** to agent **R**.

A brief description of the main plans follow.

- Plan `test` is triggered by the initial goal `!test` that starts the testing, by setting the current state to the initial state.
- Plan `move2state` upgrades the belief about the current state.
- Plan `successfulMove` is triggered by the `!type_check_message(msg(S, R, P, C))` internal goal. If the `type_check(msg(S, R, P, C), NewState)` context is satisfied, then **S** is allowed to send the message with performative **P** and content **C** to **R**. The state of the protocol changes, and `monitor` notifies **S** that the message can be sent.
- Plan `failingMoveAndProtocol` is triggered, like `successfulMove`, by the `!type_check_message(msg(S, R, P, C))` internal goal. It is used when `successfulMove` cannot be applied because its context is not verified. This means that **S** is not allowed to send message **P** with content **C** to **R**, because a dynamic type error has been detected: the message does not comply with the protocol.
- Plan `messageReceptionOK` is triggered by the reception of a tell message with `msg(S, R, P, C)` content; the message is checked against the protocol, and the progress check is activated (`!check_progress` succeeds either if a message is received before a default timeout, or if the timeout elapses, in which case `!check_progress` is activated again: `.wait(+msg(S1, R1, P1, C1), MS, Delay)` suspends the intention until `msg(S1, R1, P1, C1)` is received or `MS` milliseconds have passed, whatever happens first; `Delay` is unified to the elapsed time from the start of `.wait` until the event or timeout).

All plans whose context involves checking the current state and/or whose body involves changing it are defined as atomic ones, to avoid problems due to interleaved check-modify actions.

```

/* Initial goals */

!test.

/* Monitor's plans */

@test[atomic]
+!test : initial_state(InitialState)
      <- +current_state(InitialState).

@move2state[atomic]
+!move_to_state(NewState) : current_state(LastState)
      <- -current_state(LastState);
      +current_state(NewState).

@successfulMove[atomic]
+!type_check_message(msg(S, R, P, C)) : type_check(msg(S, R, P, C), NewState)
      <- !move_to_state(NewState);
      .print("\nMessage ", msg(S, R, P, C), "\nleads to state ", NewState, "\n");
      .send(S, tell, ok_check(msg(S, R, P, C))).

@failingMoveAndProtocol
+!type_check_message(msg(S, R, P, C)) : current_state(Current)

```

```

    <- .print("\n*** DYNAMIC TYPE-CHECKING ERROR ***\nMessage ", msg(S, R, P, C),
        "\ncannot be accepted in the current state ", Current, "\n");
        !move_to_state(failure).

@messageReceptionOK
+msg(S, R, P, C)[source(S)]: true
  <- -msg(S, R, P, C)[source(S)];
        !type_check_message(msg(S, R, P, C));
        !check_progress.

+!check_progress : timeout(MS)
  <- .wait({+msg(S1, R1, P1, C1)}, MS, Delay);
        !aux_check_progress(Delay).

+!aux_check_progress(Delay) : timeout(MS) & Delay < MS.

+!aux_check_progress(Delay) : timeout(MS) & current_state(Current) & Delay >= MS
  <- .print("\n*** WARNING ***\nNo progress for ", Delay, " milliseconds
in the current state ", Current, "\n");
        !check_progress.

```

3.2 Participants

We assume that participants interact via asynchronous exchange of messages with `tell` performatives.

To keep the implementation as general and flexible as possible, in the participants' code extended as explained below we use the `Perf` logical variable where the message performative is expected. Under the assumption that only `tell` performatives will be used, `Perf` will always be bound to the `tell` ground atom.

Only two changes are required to the code of participants:

1. `.send` is replaced by `!my_send` and
2. two plans are added for managing the interaction with the monitor.

The first plan is triggered by the `!my_send` internal goal; `my_send` has the same signature as the `.send` internal action, but, instead of sending a message with performative `Perf` and `Content` to `Receiver`, it sends a `tell` message to the monitor in the format `msg(Sender, Receiver, Perf, Content)`. When received, this message will be checked by the monitor against the global type, as explained in Section 3.1.

The second plan is triggered by the reception of the monitor's message that allows the agent to actually send `Content` to `Receiver`, by means of a message with performative `Perf`. In reaction to the reception of such a message, the agent sends the corresponding message to the expected agent.

```

/* Plans for runtime type checking */

+!my_send(Receiver, Perf, Content) : true
  <- .my_name(Sender);
        .send(monitor, tell, msg(Sender, Receiver, Perf, Content)).

+ok_check(msg(Sender, Receiver, Perf, Content))[source(monitor)] : true
  <- -ok_check(msg(Sender, Receiver, Perf, Content))[source(monitor)];
        .send(Receiver, Perf, Content).

```

3.3 Experiments

Table 1 summarizes the results of some of the experiments we carried out with the brokering protocol. The full implementation of the seller, client and broker agents, as well as the messages printed by the monitor on the console are described in [2].

- *Broker: i.o.* is the initial offer the broker makes to the client.
- *Broker: a.o.* is the lowest price the broker is willing to accept for selling oranges to the client.
- *Client: c.o.* is the client’s initial counter offer.
- *Code* is the agents code used to run the experiment.
- *Expected res.* and *Obtained res.* are the expected and obtained results.
- *Bug1*: instead of sending a counter offer upon reception of the broker’s offer, the client sends an `offer` followed by a message with unknown type.
- *Bug2*: the client autonomously starts to interact with the broker before the initial messages that the protocol enforces have been sent.
- *Bug3*: we deleted all the plans triggered by the reception of `+counter(Item, Offer) [source(Client)]` from the broker’s code, making the broker agent unable to react to a counter offer.

Our “meta-testing” of the testing mechanism was successful. We run the MAS with many other values of the initial and acceptable offers, and with other communication errors, always obtaining the expected result. The simpler protocols involving `alice` and `bob` agents described in Section 2 have been successfully tested as well.

Broker: i.o.	Broker: a.o.	Client: c.o.	Code	Expected res.	Obtained res.
11	6	3	Correct	ok	ok
8	6	2	Correct	noDeal	noDeal
8	6	5	Bug1	protocol error	protocol error
8	6	5	Bug2	protocol error	protocol error
8	6	5	Bug3	no progress	no progress

Table 1. Some of our experiments with the brokering protocol

3.4 Discussion

Alternative implementations. We opted to implement the proof-of-concept of our approach by extending the code of the existing participants rather than modifying the code of the Jason interpreter, because this was the simplest and quickest solution we could devise for developing a prototype, and easily experimenting different design choices. However, the same results could be obtained by

directly modifying the `.send` internal action by overriding the underlying agent architecture methods of Jason responsible for sending and receiving messages.

This solution would not require any modification of the code of the participants, and would allow the monitor to forward the message, when correct, directly to the recipient agent, thus reducing the number of interactions required among agents.

Another interesting solution would consist in creating a monitor agent for each agent participating to the interaction, thus avoiding the communication problems of the centralized approach where the unique monitor is required to exchange a large amount of messages with the other agents; however, this solution requires to project the global session type to end-point types (a.k.a. local types), specifying the expected behavior of each single agent involved in the interaction. Depending on the considered notion of global type, it might be non trivial to find an efficient and complete projection algorithm.

Global type transition. We have already shown that the `next` predicate is ensured to terminate on contractive global types; however, a developer may erroneously define a non contractive type for testing its system. Fortunately, there exist algorithms for automatically translating a non contractive global type into an equivalent contractive one.

Another issue concerns non deterministic global types, that is, global types where transitions are not deterministic. Consider for instance the following global type:

```
fork(seq(msg(alice,bob,tell,ping),
         seq(msg(bob,alice,tell,pong),end)),
      seq(msg(alice,bob,tell,ping),
         seq(msg(alice,bob,tell,bye),end)))
```

In this case the `next` predicate has to guess which of the two operand types must progress upon reception of the message matching with `msg(alice,bob,tell,ping)`; this means that in case of non deterministic global types the monitor may detect false positives. To avoid this problem one could determinize the type, but depending on the considered notion of global type, it would not be easy, or even possible, to devise a determinization algorithm. Alternatively, the monitor could store the whole sequence of received sending actions to allow backtracking in case of failure, thus making the testing procedure much less efficient.

Finally, it is worth mentioning that the proposed approach makes an efficient use of memory space if the initial global type does not contain loops with the `fork` constructor. In this case the space required by a global type representing an intermediate state is bounded by the size of the initial global type; since only one type at a time is kept in the belief base of the monitor, this implies a significant space optimization when the total number of all possible states is exponential w.r.t. the size of the initial global type. As already pointed out, this consideration does not apply to types with loops involving the `fork` constructor, like in the following example:

```
T = fork(seq(msg(alice,bob,tell,ping),T),
```

```
seq(msg(bob,alice,tell,pong),T)).
```

In this case the term grows at each transition step (and there are cases where the type cannot be simplified to a smaller one); however, we were not able to come up with examples of realistic protocols that require types with `fork` in a loop to be specified.

4 Related and Future Work

Our work represents a first step in two directions: extending an existing agent programming language with session types, and supporting testing of protocol conformance within a MAS. In this section we consider the related works in both areas, discuss the (lack of) proposals of integrating session types in existing MASs frameworks, and outline possible extensions of our work.

Session types on top of existing programming languages. The integration of session types into existing languages is a recent activity, dating back to less than ten years ago for object oriented calculi, and less than five years for declarative ones. The research field is very lively and open, with the newest proposals published just a few months ago.

Session types have been integrated into object calculi starting from 2005 [8, 9]. The first full implementation of a language and run-time for session-based distributed programming on top of Java, featuring asynchronous message passing, delegation, session subtyping and interleaving, combined with class downloading and failure handling, dates back to 2008 [13]. More recently, a Java language extension has been proposed, that counters the problems of traditional event-based programming with abstractions and safety guarantees based on session types [12].

Closer to our work on declarative languages, the paper [18] discusses how session types have been incorporated into Haskell as a standard library that allows the developer to statically verify the use of the communication primitives provided without an additional type checker, preprocessor or modification to the compiler. A session typing system for a featherweight Erlang calculus that encompasses the main communication abilities of the language is presented in [16]. Structured types are used to govern the interaction of Erlang processes, ensuring that their behavior is safe with respect to a defined protocol.

Protocol representation and verification in MASs. Because of the very nature of MASs as complex systems consisting of autonomous communicating entities that must adhere to a given protocol in order to allow the MAS correct functioning, the problem of how representing interaction protocols has been addressed since the dawning of research on MASs (one of the most well known outcomes being FIPA AUML interaction diagrams [14]), and the literature on protocol conformance verification is extremely rich.

Although a bit dated, [3] still represents one of the most valuable contributions to *verification of a priori conformance*. In that paper the authors propose

an approach based on the theory of formal languages. The ability to formally prove the interoperability of two policies (the actual protocol implementations), each of which is compliant with a protocol specification, is one of the main features of the proposed approach whose aim is however deeply different from ours, being devoted to a static analysis carried out before the interaction takes place.

The problem of *verifying the compliance of protocols at run time* has been tackled – among others – within the SOCS project⁷, where the SCIFF computational logic framework [1] is used to provide the semantics of social integrity constraints. Such a semantics is based on abduction: expectations on the possibly observable, yet unknown, events are modeled as abducibles and social integrity constraints are represented as integrity constraints. To model MAS interaction, expectation-based semantics specifies the links between the observed events and the expected ones. The recent paper “Modelling Interactions via Commitments and Expectations” [20] discusses that and related approaches. Although aimed at testing run-time conformance of an actual conversation with respect to a given protocol, our approach differs from the expectation-based one in many respects, including the lack of notion of expectation in the agent language, and the implementation of the testing mechanism in a seamless way on top of an existing and widespread agent-oriented programming language.

As far as formalisms for representing agent interaction protocols are concerned, the reader may find a concise but very good survey in Section 4 of [19]. In that paper, the authors propose a commitment-based semantics of protocols. Commitments involve a debtor, a creditor, an antecedent, and a consequent: the debtor stakes a claim or makes a promise to the creditor about the specified consequent provided that the antecedent holds. Protocols specify business interactions by stating how messages affect the participants’ commitments. That setting allows the authors to determine if a protocol refines another protocol, how protocols may be aggregated into other protocols, and to verify interoperability properties of agents and roles (safety, liveness, or alignment), conformance of roles, and compliance of agents. Our approach is currently limited to the runtime verification of the MAS compliance to the interaction protocol, but the exploitation of session types as the formalism to represent protocols allows us to take advantage of all the results achieved in the session types research field, which include session subtyping and algorithms for static verification of protocol properties such as safety and liveness.

The ability to specify the type of messages (`has_type(c, ct)` predicate) in order to relate actual messages to messages specified in the protocol, usually given at a more abstract level, is a characterizing feature of our approach and seems to be supported by none of the proposals mentioned above.

Session Types and MASs. As demonstrated for example by the Scribble language mentioned in the Introduction and by [10], using session types to represent and verify protocol conformance inside MASs is not a new idea but, to the best of our knowledge, no attempts of taking advantage of global session types to verify

⁷ <http://lia.deis.unibo.it/research/projects/SOCS/>

MASs programmed in some widespread agent oriented programming languages had been made so far, and our proposal is an original one.

Future extensions. Our work can be extended in many ways, as already discussed to some extent in Section 3. Besides the specific extensions mentioned there, and the fully automatic generation of the monitor and participants code, our short term goals include analyzing how our approach could be extended to other Prolog-based agent-programming languages, such as GOAL [4] or 2APL [6], and designing more complex protocols to stress-test our system and provide a quantitative assessment of its runtime behavior and scalability.

In the medium term, we plan to work for evolving our mechanism towards a framework supporting self-recovering MASs. This evolution would require to modify the way we extend the code of the participant agents, in order to automatically select other messages to send in the current state, if any, in case the monitor realizes that the chosen one does not respect the protocol. Default recovery actions for the situation where no other choices are available, should be defined as well. In such a context – more oriented towards verification of interoperability of deployed systems rather than testing of systems-to-be –, agents might advertise to the monitor the services they offer and the protocols to follow in order to obtain them. Besides ensuring the protocol’s compliance, the monitor could then act as a repository of $\langle \textit{service specification}, \textit{protocol specification} \rangle$ couples, helping agents to locate services in an open MAS in a similar way the Universal Description, Discovery and Integration (UDDI) registry does for web services.

In the long term, the integration of ontology-based meaning into protocol specifications, leading to “ontology-aware session types”, will be addressed. Our previous work on Cool-AgentSpeak [15] will represent the starting point for that extension.

Acknowledgments

We are grateful to J. F. Hübner and R. H. Bordini for their effort in making cyclic terms in Jason belief base persistent, thus making the implementation of our monitor agent possible. We also thank the anonymous reviewers for their careful reading of the paper and for the valuable suggestions provided to improve its quality.

References

1. M. Alberti, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. The SCIFF abductive proof-procedure. In *AI*IA*, pages 135–147, 2005.
2. D. Ancona, S. Drossopoulou, and V. Mascardi. Automatic generation of self-monitoring MASs from multiparty global session types - extended version. Technical report, University of Genova, Department of Computing, 2012. Online at <http://www.disi.unige.it/person/MascardiV/Download/sessionTypes4MASs.pdf>.

3. M. Baldoni, C. Baroglio, A. Martelli, and V. Patti. Verification of protocol conformance and agent interoperability. In *CLIMA VI, 2005, Revised Selected and Invited Papers*, volume 3900 of *LNCS*, pages 265–283. Springer, 2005.
4. L. Braubach, A. Pokahr, D. Moldt, and W. Lamersdorf. Goal representation for BDI agent systems. In *ProMAS 2004, Selected Revised and Invited Papers*, volume 3346 of *LNCS*, pages 44–65. Springer, 2004.
5. M. Carbone, K. Honda, and N. Yoshida. Structured communication-centred programming for web services. In *ESOP'07 (part of ETAPS 2007)*, volume 4421 of *LNCS*, pages 2–17. Springer, 2007.
6. M. Dastani. 2APL: a practical agent programming language. *Autonomous Agents and Multi-Agent Systems*, 16(3):214–248, 2008.
7. P.-M. Deniérou and N. Yoshida. Multiparty session types meet communicating automata. In *ESOP'12 (part of ETAPS 2012)*, LNCS. Springer, 2012.
8. M. Dezani-Ciancaglini, D. Mostrous, N. Yoshida, and S. Drossopoulou. Session types for object-oriented languages. In *ECOOP 2006*, volume 4067 of *LNCS*, pages 328–352. Springer, 2006.
9. M. Dezani-Ciancaglini, N. Yoshida, A. Ahern, and S. Drossopoulou. A distributed object-oriented language with session types. In *TGC 2005, Revised Selected Papers*, volume 3705 of *LNCS*, pages 299–318. Springer, 2005.
10. C. Grigore and R. Collier. Supporting agent systems in the programming language. In *WI/IAT*, pages 9–12. IEEE Computer Society, 2011.
11. K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *POPL 2008*, pages 273–284. ACM, 2008.
12. R. Hu, D. Kouzapas, O. Pernet, N. Yoshida, and K. Honda. Type-safe eventful sessions in Java. In *ECOOP 2010*, volume 6183 of *LNCS*, pages 329–353. Springer, 2010.
13. R. Hu, N. Yoshida, and K. Honda. Session-based distributed programming in Java. In *ECOOP 2008*, volume 5142 of *LNCS*, pages 516–541. Springer, 2008.
14. M.-P. Huget, B. Bauer, J. Odell, R. Levy, P. Turci, R. Cervenka, and H. Zhu. FIPA modeling: Interaction diagrams. Working Draft Version 2003-07-02. Online at <http://www.auml.org/auml/documents/ID-03-07-02.pdf>.
15. V. Mascardi, D. Ancona, R. H. Bordini, and A. Ricci. Cool-AgentSpeak: Enhancing AgentSpeak-DL agents with plan exchange and ontology services. In *IAT 2011*, pages 109–116. IEEE Computer Society, 2011.
16. D. Mostrous and V. T. Vasconcelos. Session typing for a featherweight Erlang. In *COORDINATION 2011*, volume 6721 of *LNCS*, pages 95–109. Springer, 2011.
17. A. S. Rao. AgentSpeak(L): BDI agents speak out in a logical computable language. In *MAAMAW'96*, volume 1038 of *LNCS*, pages 42–55. Springer, 1996.
18. M. Sackman and S. Eisenbach. Session types in Haskell: Updating message passing for the 21st century. Technical report, Imperial College, Department of Computing, 2008. Online at <http://spiral.imperial.ac.uk:8080/handle/10044/1/5918>.
19. M. P. Singh and A. K. Chopra. Correctness properties for multiagent systems. In *DALT 2009, Revised Selected and Invited Papers*, volume 5948 of *LNCS*, pages 192–207. Springer, 2009.
20. P. Torroni, P. Yolum, M. P. Singh, M. Alberti, F. Chesani, M. Gavanelli, E. Lamma, and P. Mello. Modelling interactions via commitments and expectations. In *Handbook of Research on Multi-Agent Systems: Semantics and Dynamics of Organizational Models*. IGI Global, 2009.