

## MULTIPARTY SESSION ACTORS\*

RUMYANA NEYKOVA AND NOBUKO YOSHIDA

Imperial College London, Department of Computing, 180 Queens Gate, London, UK  
*e-mail address:* {rn710, n.yoshida}@imperial.ac.uk

**ABSTRACT.** Actor coordination armoured with a suitable protocol description language has been a pressing problem in the actors community. We study the applicability of multiparty session type (MPST) protocols for verification of actor programs. We incorporate sessions to actors by introducing minimum additions to the model such as the notion of actor roles and protocol mailboxes. The framework uses Scribble, which is a protocol description language based on multiparty session types. Our programming model supports actor-like syntax and runtime verification mechanism guaranteeing communication safety of the participating entities. An actor can implement multiple roles in a similar way as an object can implement multiple interfaces. Multiple roles allow for cooperative inter-concurrency in a single actor. We demonstrate our framework by designing and implementing a session actor library in Python and its runtime verification mechanism. Benchmark results demonstrate that the runtime checks induce negligible overhead. We evaluate the applicability of our verification framework to specify actor interactions by implementing twelve examples from an actor benchmark suit.

### 1. INTRODUCTION

The actor model [2, 3] is (re)gaining attention in the research community and in the mainstream programming languages as a promising concurrency paradigm. Unfortunately, the programming model itself does not ensure correct sequencing of interactions between different computational processes. A study in [35] points out that the property of no shared space and asynchronous communication can make implementing coordination protocols harder and providing a language for coordinating protocols is needed. This reveals that the coordination of actors is a challenging problem when implementing, and especially when scaling up an actor system.

To overcome this problem, we need to solve several shortcomings existing in the actor programming models. First, although actors often have multiple states and complex policies for changing states, no general-purpose specification language is currently in use for describing actor protocols. Second, a clear guidance on actor discovery and coordination of distributed actors is missing. This leads to adhoc implementations and mixing the model with other paradigms which weaken its benefits [35]. Third, no verification mechanism (neither static nor dynamic) is proposed to ensure correct sequencing of interactions in a system of actors. Most actor implementations provide static typing within a single actor (even with expressive pattern matching capabilities as surveyed in § 7),

*Key words and phrases:* Session types, Protocol description language, Actors, Python, Messaging Middleware.

\* A Preliminary version of this article appeared in COORDINATION 2014.

but the communication between actors – the complex communication patterns that are most likely to deadlock – are not checked.

We tackle the aforementioned challenges by studying applicability of multiparty session types (MPST) [19], a type theory for communicating processes, to actor systems. The tenet of MPST safety assurance methodology is the use of a high-level, global specification (also called protocol) for describing interactions between communicating entities. From a global protocol each entity is given a local session type which defines the order and the payload type of interactions. Programs are then written as a collection of possibly interleaved conversations, verified, against prescribed protocols and constraints, at runtime. The practical incarnation of the theoretical MPST is the protocol specification language Scribble [32]. Scribble is equipped with a verification mechanism, which makes protocol creation easier and protocol verification sound. Declarative protocol specifications in Scribble can readily avoid typical errors in communications programming, including type errors, disrespect of call orders, circular service dependencies and deadlocks. We call these safety properties, ensured by MPST, *communication safety*.

Recent works from [25, 20, 12] prove the suitability of Scribble and its tools for dynamic verification of real world complex protocols [27] and present a runtime verification framework that guarantees safety and session fidelity of the underlying communications. The verification mechanism is applied to a large cyberinfrastructure. The MPST dynamic verification framework is built on a runtime layer for protocol management and developers use MPST primitives for communication, which limits the verification methodology to a specific infrastructure.

In this article, we take the MPST one step further. We adapt and distil the model to present a MPST verification of actors systems. A main departure from our previous works [20, 12] is that they require special conversation runtime, built into the application, which restrict the programming style and the model applicability. In this article, we prove the generality of MPST framework by showing Scribble protocols offer a wider usage, in particular, for actor programming.

Our programming model is grounded on three new design ideas: (1) use Scribble protocols and their relation to finite state machines for specification and runtime verification of actor interactions; (2) augment actor messages and their mailboxes dynamically with protocol (role) information; and (3) propose an algorithm based on virtual routers (protocol mailboxes) for the dynamic discovery of actor mailboxes within a protocol. We implement a session actor library in Python to demonstrate the applicability of the approach. To the best of our knowledge, this is the first design and implementation of session types and their dynamic verification toolchain in an actor library.

**Outline.** This article is an expansion of the initial presentation of our MPST actor framework and implementation of monitors in [24]. Apart from the expanded explanations, this article adds a new section which illustrates representative actor use cases from the literature [21, 29]. We examine expressiveness of MPSTs for specifying actor interactions by implementing protocols of use cases in an extension of Scribble; and their monitoring in Python.

The article is organised as follows: § 2 gives a brief overview of the the key features of our model and presents the running example. § 3 describes the constructs of Session Actors and highlights the main design decisions. § 4 presents the implementation of the model on concrete middleware. § 5 evaluates the overheads of the framework, compares it with a popular Scala actor library [4] and shows applications. § 6 presents and implements representative actor use cases in our framework. Finally, § 7 discusses related work and concludes. The code for the runtime and the monitor tool, example applications and omitted details are available at [33].

## 2. SESSION ACTORS PROGRAMMING MODEL

This section explains how we incorporate the verification model of MPST with the programming model of actors. The top-down development methodology for MPST verification, which relies on a standalone tool for writing protocols in the Scribble language, is shown on the left part of Fig. 1. The figure on the right illustrates the Session Actor programming model, which consists of annotated Python classes.

### 2.1. Background.

**The Actor model.** We assume the following actor features to determine our design choices. Actors are concurrent autonomous entities that exchange messages asynchronously. An actor is an entity (a class in our framework) that has a mailbox and a behaviour. Actors communicate between each other only by exchanging messages. Upon receiving a message, the behaviour of the actor is executed, upon which the actor can send messages to other actors, create actors or change its internal state. Each actor is equipped with a mailbox where messages are buffered. There are only three primitive operations each actor can perform: (1) create new actors; (2) send messages to other actors whose address are known; and (3) become an actor with a new state and behaviour. All these operations are atomic. The only order inherent is a causal order, and the only guarantee provided is that messages sent by actors will eventually be processed by their receivers. In addition, in our framework we enforce synchronisation and coordination constraints.

**Scribble overview.** As a specification language we use Scribble [32, 18], which is a practical and human-readable language for protocol specification, designed from the multiparty session types theory [19, 6]. It is a language to describe application-level protocols among communicating systems and aims to prevent communication mismatches and type errors during communications. A Scribble protocol represents an agreement on how participating systems interact with each other by describing an abstract structure of message exchanges between roles. Roles abstract from the actual identity of the endpoints that participate in a protocol session.

The Scribble tools include libraries for parsing, well-formedness checking algorithms and several language bindings. We use a Python implementation of the Scribble projection tool, available from [17]. The tool generates local Scribble specifications from global Scribble protocols. A local protocol specifies the message interactions from the viewpoint of only one of the participating entities. The Session Actor Framework retrieves the local protocols, stored in the file system or at a known location in the network, and use them for validation.

The syntax of Scribble global protocols is given by the grammar below where names are ranged over by *pro*, and the messages are of the form  $a(T)$  with  $a$  being a label and  $T$  being the constant type of the message exchanged (such as *real*, *bool* and *nat*).

$S$	::=	<code>global protocol pro (role <math>A_1, \dots, \text{role } A_n)\{G\}</math></code>	<i>specification</i>
$G$	::=	<code>a(T) from A to B; G</code>	<i>interaction</i>
		<code>choice at p {<math>G_1</math>} or ... or {<math>G_n</math>}</code>	<i>choice</i>
		<code>par {<math>G_1</math>} and ... and {<math>G_n</math>}</code>	<i>parallel</i>
		<code>rec pro {<math>G</math>}</code>	<i>recursion</i>
		<code>continue pro</code>	<i>call</i>

A (global) specification  $S$  declares a protocol with name *pro*, involving a list  $(A_1, \dots, A_n)$  of roles, and prescribing the behaviour in  $G$ . The other constructs are explained below:




FIGURE 1. MPST-development methodology (left) and session actors (right)

- An interaction,  $a(T)$  from A to B;  $G$ , specifies that a message  $a(T)$  should be sent from role A to role B and that the protocol should then continue as prescribed by the continuation  $G$ .
- A choice,  $\text{choice at A } \{G_1\} \text{ or } \dots \text{ or } \{G_n\}$  specifies a branching where role A chooses to engage in the interactions prescribed by one of the options  $G$ . The decision itself is an internal process to role A, i.e. how A decides which scenario to follow is not specified by the protocol.
- A parallel,  $\text{par } \{G_1\} \text{ and } \dots \text{ and } \{G_n\}$ , specifies unordered sequence of interactions between the specified branches  $G_1 \dots G_n$ . The interactions in  $G_1$  interleave with any of the interactions in  $G_j$  such that  $j \in (1, n]$ .
- A recursion,  $\text{rec pro } \{G\}$ , defines a scope with a name  $\text{pro}$  and a body  $G$ . Any call  $\text{continue pro}$  occurring inside  $G$  executes another recursion instance. If  $\text{continue pro}$  is not in an appropriate scope than it remains idle.

Causality in Scribble protocols is determined solely by the explicitly specified message transfers and, in particular, by the configuration of roles performing send and receive actions in message transfer sequences. Message transfers are ordered between each pair of roles. Thus, for any pair of roles A and B in a given protocol, messages sent by A to B will be received by B in the same order that they were originally dispatched by A. Consider the following partial protocol:

```
m1() from A to B;
m2() from C to A;
```

The example above requires ordering on A, but it does not specify synchronisation between B and C. Indeed, C can send the message  $m_2$  at any time, but the endpoint code for A is allowed to read it from A's mailbox only after A has sent the message  $m_1$  to B.

**MPST verification methodology.** The top-down development methodology for MPST verification is shown on the left part of Fig. 1.

A distributed protocol is specified as a global Scribble protocol, which collectively defines the admissible communication behaviours between the participating entities, called *roles*. Then, the Scribble toolchain is used to algorithmically project the global protocol to local specifications.

For each role a finite state machine (FSM) is generated from the local specification. The FSM prescribes the allowed behaviour for each role and is used at runtime to perform validation of the allowed interactions. When a party requests to start or join a session, the initial message specifies which role it implements. Its monitor retrieves the local specification based on a protocol name and a

role name. A runtime monitor ensures that each endpoint program conforms to the core protocol structure.

**2.2. Session Actors Framework.** We explain how we apply the above methodology. As observed in [31] in many actor systems an actor encapsulates a number of passive objects, accessed from the other actors through asynchronous method calls. Similarly, we introduce *multiparty roles*, which allow multiple local control flows inside an actor. We call an actor hosting *multiparty roles* a Session Actor.

**Session Actors.** To verify actor interactions, each role inside an actor is associated with a FSM, generated from a global protocol. Roles are enabled on an actor instance via two annotations, a protocol and a role annotation. The former specifies a role declaration, while the latter specifies the role behaviour. Without the introduced annotations a session actor behaves identically to a plain actor.

The right part of Fig. 1 shows the main actions performed by the framework and explained below:

- (1) `create` denotes the creation of a role inside an actor. The `@protocol` annotation creates a role object when the actor class is spawned. The role is associated to a protocol and contains proxies (remote references) for other actors (the other roles in the protocol) and a monitor.
- (2) `load` denotes loading a local protocol from the protocol storage when the role is created. The role monitor loads the protocol using the protocol name and generates a FSM that is used for checking.
- (3) `check` denotes checking of a message. When a message is picked up from the actor queue, it is checked by the role monitor and, if correct, dispatched to its handler.
- (4) `dispatch` denotes dispatching the message to the message handler. When a message arrives in the actor queue, it is dispatched to the according handler (method) according to the meta information stored in the message (a target role and a message label).

An example of how the `@role` annotation can be used is shown below. First, consider the partial protocol

```
m1() from seller to buyer;
m2() from buyer to seller;
```

Then, the message transfer declaration

```
m1() from seller to buyer;
```

corresponds to a method declaration

```
@role(seller, buyer)
def m1(self, seller, ...)
```

In a nutshell, an actor can be transformed to a session actor by applying the following design methodology:

- (1) the global protocol the actor is part of is written in Scribble and projected to local specifications using the Scribble commandline tool available from [32]. The tool generates one file (local protocol) per role.
- (2) the actor class is annotated with a `@protocol` decorator.
- (3) each method is annotated with a `@role` decorator.
- (4) interactions with other actors are performed via the role instance.

**Protocol mailbox.** Protocol mailboxes are the post offices of our systems and they exist for the duration of a protocol execution. Essentially they are virtual routers, which delegate messages according to their routing table. A protocol mailbox should be explicitly created by the developer using the `Protocol.create` primitive. This primitive generates a fresh protocol id (the name of the virtual router) and sends the id to the actors implementing the roles for the protocol. The metadata of each message that is sent from a role instance (For example, `seller.buyer.send.m2()` contains the protocol id and the target role for the message. The message is routed to the protocol mailbox matching the protocol id. The protocol mailbox delegates the message to its final recipient - the mailbox of the actor that implements the target role.

The changes and additions that the MPST annotations bring to the original actor model are as follows (here we give a high level overview, while the technical details regarding their implementation are given in subsequent sections):

- (a) different passive objects (called roles) that run inside an actor are given a session type (§ 4.2).
- (b) an actor engages in structured protocol communications via protocol mailboxes (§ 4.3).
- (c) an actor message execution is bound to a protocol structure (§ 4.4). This structure is checked via the internal FSM-based monitor.

**Verification of Session Actors.** To guarantee session fidelity (actors follow the behaviour prescribed by protocols) we perform two main checks. First, we verify that the type (a label and a payload) of each message matches its specification. In our implementation we have mapped message transfer labels to method invocations. For example,

```
m1(int) from A to B;
```

specifies that there is a method named `m1` with an argument of type `int` in the actor, implementing role `A`. Therefore, if such a method does not exist or its signature does not match, an error will be detected.

Second, we verify that the overall order of interactions is correct, i.e. interaction sequences, branches and recursions proceed as expected, respecting the explicit dependencies. Consider the following protocol:

```
m1() from A to B;
m2() from B to C;
```

The protocol specifies that there is a causality between `m1` and `m2` at the role `B`. More precisely, role `B` is not allowed to send the message `m2()` to `C` before receiving the message `m1()` from `A`. Thus, errors, such as communication mismatches, that violate the permitted protocol are ruled out.

If all messages comply to their assigned FSMs, the whole communication is guaranteed to be safe. If an actor does not comply, violations (such as deadlocks, communication and type mismatch) are detected dynamically. The monitor can either block the faulty messages (outgoing faulty messages are not dispatched to the network, incoming faulty messages are not dispatched to the message handlers) or issue a warning.

**2.3. Warehouse Management Use Case.** To illustrate and motivate central design decisions of our model, we present the buyer-seller protocol from [19] and extend it to a full warehouse management scenario. A warehouse consists of multiple customers communicating to a warehouse provider. It can be involved in a `Purchase` protocol (with customers), but can also be involved in a `StoreLoad` protocol with dealers to update its storage.

```

1 global protocol Purchase(
2   role B, role S, role A)
3 {
4   login(str) from B to S;
5   login(str) from S to A;
6   auth(str) from A to B, S;
7   choice at B
8     {req(str) from B to S;
9      (int) from S to B;
10  }
11  or
12  {buy(str) from B to S
13   deliver(str) from S to B;
14  }
15  or
16  {quit() from B to S;}}

1 global protocol StoreLoad
2   (role D, role S)
3 {
4   rec Rec{
5     choice at S
6       {req(str, int) from S to D;
7        put(str, int) from D to S;
8       continue Rec;}
9     or
10    {quit() from S to D;
11     acc() from D to S;}}}}

```

FIGURE 2. Global protocols in Scribble for (left) Purchase protocol and (right) StoreLoad protocol

**Scribble protocol.** The interactions between the entities in the system are presented as two Scribble protocols, shown in Fig. 2(left) and Fig. 2 (right). The protocols are called `Purchase` and `StoreLoad` and involve three (a buyer (B), a seller (S) and an authenticator (A)) and two (a store (S), a dealer (D)) parties, respectively. At the start of a purchase session, B sends login details to S, which delegates the request to an authentication server. After the authentication is completed, B sends a request quote `req` for a product to S and S replies with the product price. Then B has a choice to ask for another product (`req`), to proceed with buying the product (`buy`), or to quit (`quit`). By buying a product the warehouse decreases the product amount it has in the store. Products in stock are increased as prescribed by the `StoreLoad` protocol, Fig. 2 (right). The protocol starts with a recursion where a warehouse (in the role of S) has a choice to send a product request, of the form `req(string, int)`, to a dealer D. The first argument of `req` specifies the name of a requested product as a `string`, the second argument is of type `int` and stands for the number of requested products. After receiving the request, D delivers the product (`put` in line 7). These interactions are repeated in a loop until S decides to quit the protocol (line 10).

**Challenges.** There are several challenging points implementing the above scenario. First, a warehouse implementation of the Store role should be involved in both protocols, therefore it can play more than one role. Second, initially the user does not know the exact warehouse it is buying from, therefore the customer learns dynamically the address of the warehouse. Third, there can be a specific restriction on the protocol that cannot be expressed as system constants (such as specific timeout depending on the customer). The next section explains the implementation of this example in Session Actors.

### 3. SESSION ACTOR LANGUAGE

This section explains the main operations in the session actor language and the implementation of the running example in Python.

Conversation API operation	Purpose
<code>@protocol(self_role, protocol_name, self_role, other_roles)ActorClass</code>	Annotating an actor class
<code>@role(self_role, sender_role)msg_handler</code>	Annotating a message handler
<code>Protocol.create(protocol_name, mapping)</code>	Initiating a conversation
<code>c.role.send.method(payload)</code>	Sending a message to an actor in a role
<code>join(self, role)</code>	Actor handler for joining a protocol
<code>actor.send.method(payload)</code>	Sending a message to a known actor

FIGURE 3. Session Actor operations

**Session Actor language operations.** Fig. 3 presents the main session actor operations and annotations. The central concept of the session actor is the *role*. A role is a passive object inside an actor and it contains meta information used for MPST-verification. A session actor is registered for a role via the `@protocol` annotation. The annotation specifies the name of the protocol and all roles that are part of the protocol. The aforementioned meta information is stored in a role instance created in the actor. To be used for communication, a method should be annotated with a role using a `@role` decorator and passing a role instance name. The instance serves as a container for references to all protocol roles, which allows sending a message to a role in the session actor without knowing the actor location. A message is sent via `c.role.send.method`, where `c` is the self role instance and `role` is the receiver of the message.

Sending to a role without explicitly knowing the actor location is enabled through a mechanism for actor discovery (explained in § 4.3). When a session is started via the `create` method, actors receive an invitation for joining a protocol. The operation `join` is a default handler for invitation messages. If a session actor changes the join behaviour and applies additional security policies to the joiners, the method should be overloaded. Introducing actor roles via protocol executions is a novelty of our work: without roles and an actor discovery mechanism, actors would need additional configurations to introduce their addresses, and this would grow the complexity of implementations.

**Warehouse service implementation.** We explain the main constructs through an implementation of a session actor for the Warehouse service. Listing 1 presents the implementation of a warehouse service as a single session actor that keeps the inventory as a state (`self.purchaseDB`). Lines 1–2 annotate the session actor, e.g class `Warehouse`, with two protocol decorators – `c` and `c1` (for seller and store roles respectively). The instances `c` and `c1` are accessible within the warehouse actor and are holders for mailboxes of the other actors, involved in the two protocols.

All message handlers are annotated with a `@role` and are implemented as methods. For example, the `login` method (line 5) is invoked when a `login` message (line 4, Fig. 2 (left)) is sent. The role annotation for `c` (line 4) specifies the originator of the message to be `buyer`.

The handler body continues following line 5, Fig. 2 (left) – sending a `login` message via the `send` primitive to the session actor, registered as a role `auth` in the protocol of `c`. Value `c.auth` is initialised with the `auth` actor mailbox as a result of the actor discovery mechanism (explained in § 4.3). The handling of `auth` (line 6, Fig. 2 (left)) and `req` (line 6, Fig. 2 (right)) messages is similar, so we omit it and focus on the `buy` handler (line 9–12), where after sending the delivery details (line 11), the warehouse actor sends a message to itself (line 12) using the primitive `become` with value `update`. Value `update` is annotated with another role `c1`, but has as a sender `self`. This is the mechanism used for switching between roles within an actor. Update method (line 18–19) implements the request branch (line 6–8, Fig. 2 (right)) of the `StoreLoad` protocol – sending a request to the `dealer` and handling the reply via method `put`.



LISTING 1. Session Actor implementation for the Warehouse role

```

1  @protocol(c, Purchase, seller, buyer, auth)
2  @protocol(c1, StoreLoad, store, dealer)
3  class Warehouse(SessionActor):
4      @role(c, buyer)
5      def login(self, c, user):
6          c.auth.send.login(user)
7
8      @role(c, buyer)
9      def buy(self, c, product):
10         self.purchaseDB[product]-=1;
11         c.buyer.send.delivery(product.details)
12         self.become(update, product)
13     @role(c, buyer)
14     def quit(self, c):
15         c.buyer.send.acc()
16
17     @role(c1, self)
18     def update(self, c1, product):
19         c1.dealer.send.req(product, n)
20
21     @role(c1, dealer)
22     def put(self, c1, product):
23         self.purchaseDB[product]+=1:
24
25     def start_protocols():
26         c = Protocol.create(Purchase, {seller: Warehouse,
27 buyer: Customer, auth: Authenticator})
28         c1 = Protocol.create(StoreLoad, {store: c.seller,
29 dealer: Shop})
30         c.buyer.send.start()

```

The correct order of messages is verified by FSMs, embedded in *c* and *c1* respectively. As a result, errors, such as calling `put` before `update` or executing two consecutive updates, will be detected as invalid.

A protocol is started using a function `Protocol.create`. It takes, as parameters, (1) a name of a protocol and (2) a mapping between role names and actor classes. Note that, for readability, we use more descriptive role names –`buyer`, `seller` and `auth` instead of `B`, `S`, `A` respectively. Creating a protocol involves (1) creating a protocol mailbox with a newly generated protocol id; (2) spawning all actors (if they still have not been started) involved in the protocol; and (3) sending the address of the protocol mailbox (the protocol id) to all actors. Lines 26 starts the purchase protocol and spawns the Warehouse, Customer and Authenticator actors (Warehouse, Customer and Authenticator are names of actor classes). Lines 28 starts the `storeLoad` protocol. Since the Warehouse actor is already spawned we pass a reference to an existing actor, e.g `c.seller`. Then, the code for starting a role, `c.buyer.send.start()`, sends a message to `buyer` to start the protocol interactions as the global protocol prescribes (the protocol prescribes the first interaction is from `buyer` sending a `login` message to `seller`).

#### 4. IMPLEMENTATIONS OF SESSION ACTORS

This section explains our implementation of Session Actors. The key design choices follow the actor framework explained in § 2.2. We have implemented the multiparty session actors on top of Celery [9] (a Python framework for distributed task processing) with support for distributed actors [1]. Celery uses advanced message queue protocol (AMQP 0-9-1 [5]) as a transport. The reason for choosing AMQP network as base for our framework is that AMQP middleware shares a similar abstraction with the actor programming model, which makes the implementation of distributed actors more natural.

**4.1. AMQP Background.** We first summarise the key features of the AMQP model. In AMQP, messages are published by producers to entities, called exchanges (or mailboxes). Exchanges then distribute message copies to queues using binding rules. Then AMQP brokers (virtual routers) deliver messages to consumers subscribed to queues. Exchanges can be of different types, depending on the binding rules and the routing strategies they implement. We use three exchange types in our implementation: *round-robin* exchange (deliver messages alternating all subscribers), *direct* exchange (subscribers subscribe with a binding key and messages are delivered when a key stored in the message meta information matches the binding key of the subscription) and *broadcast* exchange (deliver a message to all subscribers).

Distributed actors are naturally represented in this AMQP context using the abstractions of exchanges. Each actor type is represented in the network as an exchange and is realised as a consumer subscribed to a queue based on a pattern matching on the actor id.

**4.2. Actor Roles.** Our extensions to the actor programming model are built using Python advanced abstraction capabilities: two main capabilities are *greenlets* (for realising the actors inter-concurrency) and *decorators* (for annotating actor types and methods).

A greenlet (or micro/green thread) is a light-weight cooperatively-scheduled execution unit in Python. A Python decorator is any callable Python object that is used to modify a function, method or class definition, annotated with @ symbol. A decorator is passed the original object being defined and returns a modified object, which is then bound to the name in the definition. The decorators in Python are partly inspired by Java annotations.

A key idea of actor roles is that each role runs as a micro-thread in an actor (using Python greenlet library). A role is activated when a message is received and ready to be processed. Switching between roles is done via the `become` primitive (as demonstrated in Listing 1), which is realised as sending a message to the internal queue of the actor. Roles are scheduled cooperatively. This means that at most one role can be active in a session actor at a time.

Actors are assigned roles by adding the `@protocol` decorator to the actor class declaration. The annotation

```
@protocol(seller, "Purchase", "seller", ["buyer"])
class Warehouse
```

creates a private field `seller` on the `class Warehouse`. The field has a type `ActorRole`. `ActorRole` is a custom class that encapsulates an actor role behaviour. In this case, the behaviour of the `seller` role. We recall Fig. 1, which illustrates the structure of a Session Actor. In addition to a protocol name, and an `ActorRole` instance (role A in Fig. 1), a Session Actor also contains a FSM monitor and addresses to other actors in the protocol. Thus, the other participating actors in the protocol are accessed using role names (for example, `seller.buyer` returns the address of the actor that implements `buyer`).




FIGURE 4. Organising Session Actors into protocols

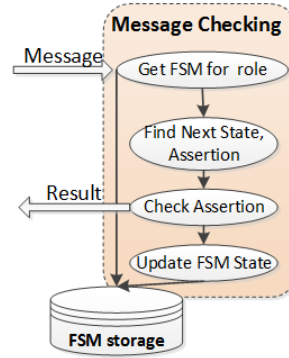


FIGURE 5. Session Actors Monitoring

All communication interactions on the role instance `seller` should follow the interactions prescribed by the protocol `Purchase` for the role `seller`. Therefore:

- (1) all labels sent to the role `seller` in the protocol specification, are implemented as methods of the class `Warehouse` and
- (2) the above methods are annotated with `@role` decorator.

For example, the specification that a buyer sends a message `quit` to the `seller` requires that there is a method `quit` annotated with `@role(seller, buyer)`. The latter can be read as `@role(receiver_role=seller, sender_role=buyer)`.

Messages received in an actor mailbox are dispatched to the matching handlers (methods) only if the `@role` decorator for the method corresponds to the meta information about the sender and receiver roles in the message header. For example, the message dispatch as a result of the call

```
buyer.seller.send.quit()
```

will invoke the body of the method declared as:

```
@role(seller, buyer)def quit.
```

**4.3. Actors Discovery.** Fig. 4 presents the network setting (in terms of AMQP objects) for realising the actor discovery for `buyer` and `seller` of the protocol `Purchase`. We use the three types from AMQP explained in § 4.1. For simplicity, we create some of the objects on starting of the actor system – round-robin exchange per actor type (`warehouse` and `customer` in Fig. 4) and broadcast exchange per protocol (`purchase` in Fig. 4). All spawned actors alternate to receive messages addressed to their exchange. When session actors are started their exchange is bound to the protocol exchange according to their `@protocol` annotations (line 1 in Listing 1 binds `warehouse` exchange to `purchase` exchange, as shown in Fig. 4).

An exchange accepts messages from a producer application and routes them to message queues. Exchanges are the “mailboxes” of AMQP. Unlike other popular middlewares, in AMQP messages

are not published directly to queues. Messages are published to exchanges that route them to queue(s) using pre-arranged criteria called bindings.

Exchanges are very cheap to create with a throughput of one million messages per second<sup>1</sup>. Therefore we do not consider their use as a bottleneck of our implementation. Messages are not stored in exchanges and therefore only the exchange throughput evaluation is relevant to performance concerns.

The session initiation takes advantage of the available AMQP infrastructure. The actor discovery wraps the boilerplate code for spawning new actors, creating new session id and exchanging actor addresses between participants. The exact exchange bindings for our running example are explained below.

We now explain the workflow for actor discovery. When a protocol is started, a fresh protocol id and an exchange with that id are created. An AMQP direct exchange (see § 4.1) is used so that messages with a routing key are delivered to actors linked to the exchange with binding to that key (it corresponds to `protocol_id` in Fig. 4). Then, a `join` message is sent to the protocol exchange and delivered to one actor per registered role (`join` is broadcasted to `warehouse` and `customer` in Fig. 4). On `join`, an actor binds itself to the `protocol_id` exchange with subscription key equal to its role (bindings `seller` and `buyer` in Fig. 4). When an actor sends a message to another actor within the same session (for example `c.buyer.send.ack()` in Listing 1), the message is sent to the protocol id exchange (stored in `c`) and from there delivered to the buyer actor.

**4.4. Preservation through FSM Checking.** Before a message is dispatched to its message handler, the message goes through a monitor. Fig. 5 illustrates the monitoring process. Each message contains meta information (a routing key) with a role name and a protocol id. When an actor joins a protocol, the FSM, generated from the Scribble compiler (as shown in Fig. 1) is loaded from a distributed storage to the actor memory.


Message checking goes through the following steps. First, depending on a role and a protocol id a matching FSM is retrieved from the actor memory. Next, the FSM checks the message labels/operators (that are part of the actor payload) and sender and receiver roles (that are part of the message binding key) are valid.

The *check assertions* step verifies that if any constraints on the size/value of the payload are specified in Scribble, they are also fulfilled. If a message is detected as wrong the session actor throws a distributed exception and sends an error message back to the sending role and does not pass the message to its handler for processing. This behaviour can be changed by implementing the `wrong_message` method of the session actor.

## 5. EVALUATIONS OF SESSION ACTORS

This section reports on the performance of our framework. The goal of our evaluation is two fold. First, we compare our host distributed actor framework [9] with a mainstream actor library (AKKA [4]). Second, we show that our main contribution, runtime verification of Scribble protocols, can be realised with reasonable cost. The full source code of the benchmark protocols and applications and the raw data are available from the project page [33].

<sup>1</sup><http://www.rabbitmq.com/blog/2011/03/28/very-fast-and-scalable-topic-routing-part-2/>



```

global protocol FlatPingPong(
  role C, role S)
{
  ping(str) from C to S;
  pong(str) from S to C;
}

global protocol RecPingPong(
  role C, role S)
{
  rec Loop{
    ping(str) from C to S;
    pong(str) from S to C;
    continue Loop; }
}

```

FIGURE 6. The PingPong benchmark, showing the overhead of message delivery (left) and The FlatPingPong and RecPingPong protocol (right)

**5.1. Session Actors Performance.** We test the overhead of message delivery in our implementation using the pingpong benchmark [22] in which two processes send each other messages back and forth. The original version of the code was obtained from Scala `pingpong.scala`<sup>2</sup> and adapted to use distributed AKKA actors (instead of local). We distinguish two protocols. Each pingpong can be a separate session (protocol `FlatPingPong`) or the whole iteration can be part of one recursive protocol (`RecPingPong`). The protocols are given in Fig. 6 (right). This distinction is important only to session actors, because the protocol shape has implications on checking. For AKKA actors the notion of session does not exist and therefore the two protocols have the same implementation.

**Set up.** We run each scenario 50 times and measured the overall execution time (the difference between successive runs was found to be negligible). The creation and population of the network was not measured as part of the execution time. The client and server actor nodes and the AMQP broker were each run on separate machines (Intel Core2 Duo 2.80 GHz, 4 GB memory, 64-bit Ubuntu 11.04, kernel 2.6.38). All hosts are interconnected through Gigabit Ethernet and latency between each node was measured to be 0.24 ms on average (ping 64 bytes). The version of Python used was 2.7, of Scala – 2.10.2 and of the AKKA-actors – 2.2.1.

**Results.** Fig. 6 (left) compares three separate benchmark measurements for session actors. The base case for comparison, annotated as "Celery", is a pure actor implementation without the addition of roles. "With Role" measures the overhead of roles annotations without having the monitor capabilities enabled (without verifying communication safety). The two main cases, "Rec Monitor" and "Monitor", measure the full overhead of session actors. This separation aims to clearly illustrate the overhead introduced by each of the additions, presented in this chapter: roles annotations, Scribble-based runtime verification and actor discovery. Note that FSMs for the recursive and for the flat protocol have the same number of states. Therefore, the observed difference in the performance is a result of the cost of actor discovery.

<sup>2</sup><http://scala-lang.org/old/node/54>

States	100	1000	10000
Time (ms)	0.0479	0.0501	0.0569

TABLE 1. Execution Time for checking protocol with increasing length

The difference between the performance of AKKA and Celery is not surprising and can be explained by the distinct nature of the frameworks. AKKA runs on top of JVM and the remote actors in these benchmarks use direct TCP connections. On the other hand, Celery is a Python-based framework which uses a message middleware as a transport, which adds a layer of indirection. Given these differences, Celery actors have reasonable performance and are a viable alternative.

Regarding our additions we can draw several positive conclusions from the benchmarks: (1) the cost of FSM checking is negligible and largely overshadowed by the communication cost (such as latency and routing); and (2) the cost of the mechanism for actor discovery is reasonable, given the protocol load.

**5.2. Monitoring Overhead.** In this subsection, we discuss the row overhead of monitor checking, which is a main factor that can affect the performance and scalability of our implementation. Knowing the complexity of checking protocols of different length is useful to reason about the applicability of the model.

We have applied two optimisations to the monitor, presented in § 4.4, which significantly reduce the runtime overhead. First, we pre-generate the FSM based on the global protocol. Second, we cache the FSM protocols the first time when a role is loaded. Therefore, the slight runtime overhead, observed in the previous section is a result of the time required to choose the correct FSM among all in-memory FSMs inside the actor and the time required to check the FSM transition table that the message is expected. Both have linear complexity on the number of roles running inside an actor and the number of states inside a FSM respectively. Table 1 shows the approximate execution time for checking a protocols of increasing length of interactions. The protocol for testing is the request-reply protocol, given below:

```
msg1() from A to B; msg2() from B to A; ... msg_n() from B to A)
```

We evaluate the above protocol for  $n$  ranging over 100, 1000 and 10 000 respectively. Then we report the time taken to complete each of the protocols. We omit results for  $n$  smaller than 100 since the execution time is negligible.

**5.3. Applications of Session Actors.** As a practical evaluation of our framework, we have implemented two popular actor applications and adapted them to session actors.

The first application is a distributed chat with multiple chat rooms, where a number of clients connect to a reactive server and execute operations based on their privileges. Depending on privileges some clients might have extended number of allowed messages. We impose this restriction dynamically by annotating the restricted operations with different roles. An operation is allowed only if its actor has already joined a session in the annotated role.

The second use case is a `MapReduce` protocol for counting words in a document. It is an adaptation from an example presented in the official website for celery actors<sup>3</sup>. The session actor usage removes the requirement for the manual actor spawning of `Reducers` actors inside the `Mapper` actor, which reduces the code size and the complexity of the implementation.

<sup>3</sup><http://cell.readthedocs.org/en/latest/getting-started/index.html>

LISTING 2. The WordCount protocol and its implementation in Python

```

1  #####Code For The Mapper Role#####
2  @protocol(c, WordCount, R, A, M)
3  class Mapper(Actor):
4      @role(c) # invoked on protocol start
5      def join(self, c, file, n):
6          self.count_document(self, file, n)
7
8      @role(c, self)
9      def count_document(self, c, file, n):
10         with open(file) as f:
11             lines = f.readlines(), count = 0
12             for line in lines:
13                 reducer = c.R[count%n], count+=1
14                 reducer.send.count_lines(line)
15
16 #####Code For Creating a protocol#####
17 c = Protocol.create(WordCount, {M=Mapper,
18     A=Aggregator}, {R:(Reducer, 10)})
19 c.M.send.join(file="file1.txt", n=10)
20
21 #####Code For Aggregator role#####
22 @protocol(c, WordCount, A, M, R)
23 class Aggregator(Actor):
24     @role(c, master)
25     def aggregate(self, c, words):
26         for word, n in words.iteritems():
27             self.result.setdefault(word, 0)
28             self.result[word] += n
29             # when a treshhold is reached
30             # print the results
31         c.close()
32
33 #####Code For Reducer role#####
34 @protocol(c, WordCount, R, A, M)
35 class Reducer(Actor):
36     @role(c, M)
37     def count_lines(self, c, line):
38         words = {}
39         for word in line.split(" "):
40             words.setdefault(word, 0)
41             words[word] += 1
42         c.A.send.aggregate(words)

```

We give in Listing 2 the implementation of the latter example. For sake of space and clarity, the implementation reported here abstracts from technical details that are not important for the scope of this work. The protocol is started on line 17 specifying actors for each role and passing the arguments for the initial execution – `n` (the number of reducers) and `file` (the name of the source file to be counted). The Mapper implements the `join` method, which is invoked when an actor joins a protocol. Since the Mapper is the starting role, when it joins, it sends a message to itself to proceed with the

No	Use Case Name	MPST feature required	paper
1	Ping Pong	standard	[18]
2	Counting Actor	standard	[25]
3	Fibonacci	recursive subsession with dynamic role creation	[11]
4	Big	parallel; join	[18]
5	Thread Ring	Pabble (parameterised MPST)	[26]
6	KFork (throughput)	subsessions	[11]
7	Dining Philosophers	parameterised parallel; group role	new
8	Sleeping Barber	parameterised parallel; group role	new
9	Logistic Map Series	parameterised parallel with join	new
10	Bank Transaction	parallel composition with join	new
11	Cigarette Smoker	parameterised parallel; group role	new

TABLE 2. List of implemented Savina benchmarks

execution of `count_document`. Note that spawning and linking of actors are not specified in the code because the actor discovery mechanism accounts for it. The protocol proceeds by `Mapper` sending one message for each line of the document. The receiver of each message is one of the `Reducers`. Each `Reducer` counts the words in its line and sends the result to the `Aggregator` (c.A, line 42), which stores all words in a dictionary and aggregates them. When a threshold for the result is reached, the `Aggregator` prints the result and stops the session explicitly.

Our experiences with session actors are promising. Although they introduce new notions, i.e. a protocol and a role, we found that their addition to a class-based actor framework is natural to integrate without requiring a radical new way of thinking. The protocol and role annotations are matched well with typical actor applications, and they even result in simplifying the code by removing the boilerplate for actor discovery and coordination. The protocol-oriented approach to actor programming accounts for the early error detection in the application design and the coordination of actors. The runtime verification guarantees and enforces the correct order of interactions, which is normally ensured only by hours of testing.

## 6. USE CASES

In this section we evaluate the applicability of our framework to verify use cases which contain common communication patterns found in actor languages. We select the use cases from an actor benchmark suit, Savina [21]. Savina aims to identify a representative set of actor applications. The use cases range from popular micro benchmarks to classic concurrency problems. Each benchmark tests a different aspect of the actor implementation (e.g. actor creation, throughput, mailbox contention) and presents a different interaction patterns (e.g. master-slave, round-robin communication, synchronous request-reply).

We express twelve out of the sixteen benchmarks, presented in the micro and concurrency benchmarks sections of the suit. We observe that standard primitives for multiparty session types (recursion, choice, point-to-point interaction and parallel composition) are not sufficient to specify most of the use cases, but more advanced features such as assertions [7], subsessions [11] and dynamic role creation [13] leverage the expressive capabilities of Scribble. The required extensions are already proposed in the theoretical literature of MPSTs, and can be expressed using combinations of various theoretical formalisms. Here we (1) identify (a combination of) the MPST extensions needed to express actors patterns; (2) specify them in Scribble; and (3) implement end-point programs in Python.



First, we explain the advanced Scribble features, pointed in Table 2. In the table, the *param-eterised parallel composition* stands for the use of the construct `par [i:1..N] {body}`, which is a shorthand for  $N$  parallel `body` branches, where  $i$  is substituted with values from 1 to  $N$  inside the `body`. When the `par` is followed by another block as in Listing 5, we call this block a join block (denoted as *parallel with join* in Table 2). The sequence `par {block 1} and {block 2} block 3` means that for `block3` to be executed, `block1` and `block2` should be completed. For example, in Listing 5 the message transfer `done() from S to A;` can be completed only after both `A` and `B` have sent a message `done` to `S`.

Another Scribble feature exploited in our examples is a *subsession*. It is used to create modular Scribble code and is realised as a macro expansion. It is often combined with a *group role* (Table 2, examples 7, 8, 11). A group role, for example, `role Philosophers[1..N]`, specifies that there will be  $N$  roles with the same type in the protocol. The roles share the same implementation, but at runtime there will be  $N$  instances of that role and their behaviour might depend on initially assigned indices.

Next we present the examples. We first describe a global Scribble protocol for each use case based on specifications in [21] and the implementations on github [29], and then provide a couple of implementations in Python in Appendix A. We also demonstrate faulty executions, that can be detected by our framework, on the Sleeping barber example in Fig. 7. The use cases in Scribble and corresponding Python code can be found in [33].

Table 2 shows a summary, specifying the communication pattern and the MPST extension (if needed). For each use case we give a brief description and the corresponding global Scribble protocol. To express the use cases the previous version of Scribble has to be extended with (1) parametrised parallel composition and (2) subsession with dynamic spawning of roles [11].

LISTING 3. the Counting Protocol

```
global protocol Counter (role P, role C)
{
  rec Loop { value(int) from P to C;}
  result(int) from C to P;}
}
```

**Counting Actor.** This use case expresses an elaborated request-reply pattern, measuring the time taken to send  $n$  consecutive messages from a `Producer` actor to a `Counter` actor. The accumulated value in the `Counter` is returned to the `Producer`. From an actor perspective, the use case tests the mail box implementation which needs to store the pending messages yet to be processed. Using session types the communication interaction is succinctly specified with recursion and point-to-point communication, expressing the causalities between sending of a value (`value(int)`) and the final reply (`result(int)`).

LISTING 4. The Fibonacci Protocol

```
global protocol Fib(role P, role C)
{
  request(int) from P to C;
  par { Fib(C as P, new C);}
  and { Fib(C as P, new C);}
  result(int) from C to P; }}
}
```

**Fibonacci.** This use case involves the recursive creation of a predefined number of actors using the naive recursive formula for computing fibonacci numbers. The idea of the protocol is that an actor receives a request to compute a value, then it spawns two additional actors and waits for their response. Finally, it returns the sum of the responses to the requester. The example tests the actor creation/destruction in the actor implementation. The global Scribble protocol is given in Listing 4. After receiving a request (`request`) from a `Parent` actor, a `Child` actor invokes the same protocol type (`Fib`) with a new `child` actor. To express the new role creation, we use nested protocols (hereafter called subprotocols), formally presented in [11]. A subprotocol definition accepts either roles already active in the protocol (`C as Parent`) or new roles of a known type (`new C`). To express the fork-join behaviour required by the example we use parallel composition with continuation (the actions after a `par` in a protocol body). Thus the specified protocol imposes causality between sending the response (`result`) and the completion of the proceeding interactions, but no causality between the subprotocols in the parallel branches.

LISTING 5. The Big protocol

```
global protocol Big (role A,
  role B, role Sink as S)
{
  par { ping() from A to B;
        pong() from B to A;
        done() from A to S;}
    and { ping() from B to A;
          pong() from A to B;
          done() from B to S;}
}
```

**Big.** The use case presents a many-to-many message passing scenario. Several actors are spawned, each of which sends a `ping` message to another random actor. Then the recipient actor responds with a `pong` message. A `Sink` actor is notified (via `done` message) about the completion of these interactions. When the `Sink` receives all replies, it terminates the other actors. In an actor implementation, the benchmark measures the effects on contention on a recipient's actor mailbox. In session types the use case requires a join primitive, expressed in Scribble as a continuation after the parallel branches, similarly to the Fibonacci use case. Note that the round robin sending cannot be precisely expressed in Scribble. We have simplified the use case, giving only the type for two ping-pong actors (`A` and `B` in Listing 5) and one `Sink` actor. The protocol can be extended for more ping-pong actors if the configuration is fixed in advance.

LISTING 6. The ThreadRing protocol

```
global protocol Ring(role Worker [1..N])
{
  rec Loop {
    data(int) from Worker[i:1..N-1] to Worker[i+1];
    data(int) from Worker[N] to Worker[1];
    continue Loop;}
}
```

**Thread Ring.** In the ThreadRing use case an integer token message is passed around a ring of `N` connected actors. The token value is decremented each time an actor receives it. The interactions continue until the token becomes zero. From an actor perspective, the benchmark tests the raw efficiency of message sending since it removes all contention of resources except between the sending

and receiving actors. To express the protocol for each  $N$ -ring configuration we use an extension of Scribble, called Pabble [26], where multiple participants can be grouped in the same role and indexed. Indeed, the same use case is already presented in [26]. In the protocol in Listing 6, `Worker[i:1..N-1]` declares iteration from 1 to  $N-1$  on the interaction (`Data(int) from Worker[i] to Worker[i+1]`) over the bound variable  $i$ .

LISTING 7. The KForK protocol

```
global protocol KForK(role master as M, role Worker[1..K])
{
  par[i:1..K] {
    rec Loop {
      choice at M {
        data(int) from M to Worker[i];}
      or {
        end from M to Worker[i];}}
  }
}
```

**KForK.** The application first creates  $k$  worker actors and then sends each actor a total of  $n$  messages in a round robin manner. Actors terminate after processing all messages. The benchmark tests the message throughput in the implementation. The global multiparty protocol is expressed as a parallel composition of  $k$  binary protocols with an identical communication pattern, recursive interaction between a `Master` and a `Worker`. For more concise specification, since all parallel branches describe an identical communication pattern, we leverage the parallel construct by introducing indexing on parallel branches. The syntax `par[i:1..N] {body}` defines replication of `body` for each value of  $i$  between 1 and  $k$ . Listing 8 shows the global protocol for KForK using the new syntax.

LISTING 8. Refactored KForK Protocol

```
global protocol KForK(role Master as M, role Worker[1..K])
{
  par[i:1..K] { Sub(M, W[i]);}
}
global protocol Sub(role M, role W)
{
  rec Loop {
    choice at M {
      data(int) from M to W;}
    or {
      end from M to W;}}
}
```

**Session patterns for concurrent group execution (Refactoring KForK).** We refactor the KForK global protocol, extracting the `par` body into a subprotocol. The global protocol with a group role is represented as a parallel composition of subprotocols. Indeed, this is an emerging pattern in all concurrent problems surveyed in the benchmark suit. The pattern appears in four use cases: Sleeping Barber, LogisticMap, BankTransaction and DiningPhilosophers.

A global protocol with a group role is split into smaller multiparty protocols - one for each role in the group role set. In addition, in LogisticMap and SleepingBarber a join multicast is required to signal the completion of the interaction and to end the protocol. In Listing 9 we give the four global protocols, where the recurring session pattern can be clearly observed. In the rest of the section we

LISTING 9. Dining Philosophers, Sleeping Barber, Bank Transaction, Logistic Map

```

global protocol DiningPhilosophers(
  role arbitrator as A, role Philosopher [1..N])
{
  par [i:1..n] { DiningPhilosopher(A, Philosopher[i]); }
}
global protocol BankTransactions(
  role Teller as T, role Accounts [1..N])
{
  par [i:1..n-1] { BankTransaction(T, A[i], A[i+1]); }
}
global protocol LogisticMap(
  role Master as M, role Series [1..N], role Ratio [1..N])
{
  request from M to S [1..N];
  par [i:1..N] { SeriesMaster(S[i], R[i]); }
  response from S [1..N] to M;
  stop from M to S [1..N];
}

```

focus on the interactions between one of the roles from the group role set and give the protocols (DiningPhilosopher, SleepingBarberOneCustomer, BankTransaction, SeriesMaster respectively) which do not require the use of a group role.

LISTING 10. The DiningPhilosophers protocol

```

global protocol DiningPhilosopher(
  role Arbitrator as A, role Philosopher as Ph)
{
  rec L { choice at Ph {
    rec M {
      req() from Ph to A;
      choice at A {
        yes() from A to Ph;
        continue L;}
      or {
        no() from A to Ph;
        continue M;}}
    or { done() from Ph to A;}}
}

```

**The Dining Philosophers.** In this classic example,  $N$  philosophers are competing to eat while sitting around a round table. There is one fork between each two philosophers and a philosopher needs two forks to eat. The challenge is to avoid deadlock, where no one can eat, because everyone is possessing exactly one fork. In Savina, the problem is implemented using `Arbitrator` actor that coordinates the resources (the forks) between `Philosopher` actors. The `Arbitrator` either accepts (`yes()`) a `Philosopher` request or rejects (`no()`) it, depending on the current resources allocated. If rejected, the `Philosopher` keeps sending request (`req`) until the request is accepted. The global

protocol in Listing 10 specifies the interactions between an `Arbitrator` and one `Philosopher`, defining the causalities using recursion and choice.

**Sleeping Barber.** The use case presents another classic synchronisation problem. A barber is waiting for customers. When there are no customers the barber sleeps. When a customer arrives, he wakes the barber or sits in one of the waiting chairs in the waiting room. If all chairs are occupied, the customer leaves. Customers repeatedly visit the barber shop until they get a haircut. The key element of the solution is to make sure that whenever a customer or a barber checks the state of the waiting room, they always see a valid state. The problem is implemented using a `Selector` actor that decides which is the next customer, `Customer` actors, a `Barber` actor and a `Room` actor. Note that, differently from the previous use cases, the protocol between the `Customers` and the other actors is multiparty (not binary). To guarantee a valid synchronisation, precise message sequence should be followed. While the use cases so far demonstrated succinct recursion and request-reply patterns, interleaved together, this one demonstrates a communication structure with long sequence of interactions, where preserving the causalities is the key to achieving deadlock free communication.

The implementation is  $\sim 200$  LOC in Akka ( $\sim 100$  LOC in Python), where each of the four actors send messages to the other three actors. The implementation is error-prone, because the type of the actors and whether they can be safely compose is not obvious. Sending the wrong message type, sending to the wrong actor or not sending in the correct message sequence may lead to deadlocks, errors which initial cause is hard to be identified, or wrong computation results. Errors can even go unnoticed resulting in orphan actors. Here we give examples of several faulty executions and explain how they are prevented by our framework (for every error example, we highlight in Fig. 7 a wrong implementation of a message transfer and the corresponding line in the global protocol):

**Deadlock:** (line 10): After receiving message `full` the `Customer` `c` sends a message `returned` to the `Room` `r` instead of sending it to the `Selector` `s`. Then the `c` will wait to be reentered again, while `s` will wait to receive `returned` or `done` from `c` (according to Line 17), resulting in a deadlock. Moreover, the other actors are stuck as well (`r` is waiting on `s`; `b` is waiting on `r`). Our framework will detect that the message `returned` is sent to the wrong actor.

**Unexpected termination:** (line 17): In this scenario, `c` again sends a wrong message to `r`, but contrary to the above example, it sends a message that `r` can handle (the message `done`). If there is no verification at `r` on the sender of the messages, `r` will execute the `done` handler and as a result will terminate (`done` should be sent by `s` to terminate the actor `r` when all customers have been processed). When the `Barber` `b` tries to send `next` message to `r` it will get a runtime exception that the actor `r` does not exist. Such an exception can be very misleading because it points to a problem between the `Barber` and the `Room`, while the actual error is in the code for the `Customer`. Our framework will detect this inconsistency at the right (at role `c`) place therefore will not allow the error to be propagated, masking the real cause.

**Orphan actors:** (line 24): The actor `r` sends a wrong message (`don` instead of `done`) to `b` (or doesn't send a message at all). Therefore, `b` will not terminate. However, everybody else will and so will the protocol, leaving the orphan actor `b`. Our monitor will detect the wrong message that `b` is sending.

**Cigarette Smoker.** This problem involves  $N$  smokers and one arbiter. The arbiter puts resources on the table and smokers pick them. A smoker needs to collect  $k$  resources to start smoking. The challenge is to avoid deadlock by disallowing a competition between smokers from picking up resources. This is done by delegating the control to the arbiter, who decides (in a random manner) which smoker to send the resource to. The session type for the use case, given in Listing 11, combines

```

1 global protocol SleepingBarber(role barber as B,
2 role selector as S, role room as R)
3 {rec L {
4     choice at S {                               # Deadlock
5     par {                                       class CustomerActor...
6     S introduces C;}                           @role(c, S)
7     enter() from S to R;                       def full(self, c):
8     choice at R {                               # c.S.send.returned()
9         full() from R to C;                     c.R.send.returned()
10    returned() from C to S;
11    continue L; }
12 } or {
13 wait() from R to C;                             ...
14     enter() from R to B;                       @role(c, B)
15     start() from B to C;                       def stop(self, c):
16     stop() from B to C;                       # c.S.send.done()
17    done() from C to S;                         c.R.send.done()
18     next() from B to R;
19 }                                               # Orphan Actors
20 wait() from R to B;                             class RoomActor...
21 continue L;                                     @role(c, S)
22 } or {                                         def done(self, c):
23     done() from S to R;                       # c.B.send.done()
24    done() from R to B; }}}}                   c.B.send.don()

```

FIGURE 7. The Sleeping Barber protocol (left) and corresponding faulty implementations (right)

round-robin pattern, sending a random smoker a message to smoke (`StartSmoking`), with multicast, iterating through all the smokers notifying them to exit (`exit from A to S[i..n]`);).

LISTING 11. The CigaretteSmoker protocol

```

global protocol CigaretteSmoker(role Arbiter as A,
role Smoker[1..N] as S)
{
rec Loop {
choice at A {
start_smoking() from A to S[i];
started_smoking() from S[i] to A;
continue Loop;
} or {
exit() from A to S[i..n]; }}}

```

**Bank Transaction.** The communication pattern in this use case is a synchronous request-reply chain. Three actors (a Teller, a source account `SrcAcc` and a destination account `DestAcc`) synchronise to perform a debit from one account to another. The Teller is not allowed other operations until a confirmation that the transaction is complete is received from the source account. The two synchronous request-reply patterns are expressed as a message sequence in the BankTransaction protocol in Listing 12.

LISTING 12. The BankTransaction protocol

```

global protocol BankTransaction(role Teller as T,
    role SrcAcc as S, role DestAcc as D)
{
    credit from T to S;
    debit from S to D;
    reply from D to S;
    reply from S to T;
}

```

**Logistic Map Series.** This is another example of a synchronous request-reply pattern. The goal is to calculate a recurrence relation:  $x_{n+1} = rx_n(1 - x_n)$ . The implementation requires an actor for a manager (Master), a set of term actors (Series) and a set of ratio actors (Ratio). The Ratio actors encapsulate the ratio  $r$  and compute the next term  $x_{n+1}$  from the current term  $x_n$ . The Master sends request to the Series actor to compute a result. Then Series actors delegate the calculation to a Ratio actor and wait for synchronous reply to update their value of  $x$ . Thus, the MasterSeries subprotocol is an instance of a request-reply pattern (from Series as a requester to Ratio role as a replier), which we have already shows and thus we omit.

LISTING 13. The ProducerConsumer protocol

```

global protocol PCBounded(
    role Buffer as B,
    role Producer as P,
    role Consumer as C)
{
    rec Loop {
        produce() from B to P;
        par {
            choice at P {
                dm() from P to B;
                continue Loop;
            } or {
                exit() from P to B;
            }
        } and {
            dm() from B to C;
            more() from C to B;
        }
    }
}

```

**Producer Consumer with Bounded Buffer.** The overall scenario of this use case is that producers push work into buffers, while consumers pull from buffers. The implementation uses actors for producers, consumers and a buffer. The Buffer actor is concurrently receiving messages from both Producers and Consumers. Thus, the benchmark measures also the contention overhead of the actor implementation. The Scribble protocol is given in Listing 13. To represent the pattern (producers are sending messages concurrently, but there is causality between the message send from a Producer and the message send to the Consumer). We use a combination of recursion and parallel branches. In each recursive execution, we require Consumers and Producers to send or receive concurrently. At the same time, the Producer can keep sending (starting a new recursive execution), before the Consumer has finished their branch.

## 7. RELATED WORK

**Behavioural and session types for actors and objects.** There are several theoretical works that have studied the behavioural types for verifying actors [23, 10]. The work [10] proposes a behavioural

typing system for an actor calculus where a type describes a sequence of inputs and outputs performed by the actor body. In [23], a concurrent fragment of Erlang is enriched with sessions and session types. Messages are linked to a session via correlation sets (explicit identifiers, contained in the message), and clients create the required references and send them in the service invocation message. The developed typing system guarantees that all within-session messages have a chance of being received. The formalism is based on only binary sessions.

Several recent papers have combined session types, as specification of protocols on communicating channels, with object-oriented paradigm. A work close to ours is [15], where a session channel is stored as a field of an object, therefore channels can be accessed from different methods. They explicitly specify the (partial) type for each method bound to a channel. Our implementation also allows modularised sessions based on channel-based communication. Their work [15] is mainly theoretical and gives a static typing system based on binary session types. Our work aims to provide a design and implementation of runtime verification based on multiparty session types, and is integrated with existing development frameworks based on Celery [9] and AMQP [5].

The work in [8] formalises behaviours of non-uniform active objects where the set of available methods may change dynamically. It uses the approach based on spatial logic for a fine grained access control of resources. Method availability in [8] depends on the state of the object in a similar way as ours.

See [37] for more general comparisons between session types and other frameworks in object-oriented paradigms.

**Other Actor frameworks.** The most popular actor’s library (the AKKA framework [4] in Scala studied in [35]) supports FSM verification mechanism (through inheritance) and typed channels. Their channel typing is simple so that it cannot capture structures of communications such as sequencing, branching or recursions. These structures ensured by session types are the key element for guaranteeing deadlock freedom between multiple actors. In addition, in [4], channels and FSMs are unrelated and cannot be intermixed; on the other hand, in our approach, we rely on external specifications based on the choreography (MPST) and the FSMs usage is internalised (i.e. FSMs are automatically generated from a global type), therefore it does not affect program structures.

Several works study an extension of actors in multicore control flows. Multithreaded active objects [16] allow several threads to co-exist inside an active object and provide an annotation system to control their concurrent executions. Parallel Actor Monitors (PAM) [31] is another framework for intra actor parallelism. PAMs are monitors attached to each actor that schedule the execution of the messages based on synchronisation constraints. Our framework also enables multiple control flows inside an actor as [16, 31]. In addition, we embed monitors inside actors in a similar way as [31, 16] embed schedulers. The main focus of the monitors in [16, 31] is scheduling the order of the method executions in order to optimise the actor performance on multi-core machines, while our approach aims to provide explicit protocol specifications and verification among interactions between distributed actors in multi-node environments.

The work [30] proposes a technique to have multiple cooperatively scheduled tasks within a single active object similar to our notion of cooperative roles within an actor. The approach is implemented as a Java extension with an actor-like concurrency where communication is based on asynchronous method calls with objects as targets. They resort to RMI for writing distributed implementations and do not allow specifying sequencing of messages (protocols) like ours.

The work [28] proposes a framework of three layers for actor roles and coordinators, which resembles roles and protocol mailboxes in our setting. Their specifications focus on QoS requirements, while our aim is to describe and ensure correct patterns of interactions (message passing).



Comparing with the above works, our aim is to provide effective framework for multi-node environments where actors can be distributed transparently to different machines and/or cores.

Several extensions of the actor model suggest more complex synchronization of the actor inbox such as multi-headed messages with guards [34], join-style actors [14] and synchronous replies [36]. These works focus on more expressive pattern matching for precise and correct dispatch of the messages to the corresponding message handlers implemented in an actor. Thus, the synchronisation constraints are specified on the local actor behaviour. The emphasize on our work is to express the synchronization constraints globally, as part of the overall choreography of an actor system, and infer (through projection) the local behaviour of a single actor.

## 8. CONCLUSION

We propose an actor specification and verification framework based on multiparty session types, providing a Python library with actor communication primitives. Cooperative multitasking within sessions allows for combining active and reactive behaviours in a simple and type-safe way. The framework is naturally integrated with Celery [9] which uses advanced message queue protocol (AMQP [5]), and uses effectively its types for realising the key mechanisms such as the actor discovery. We demonstrated that the overhead of our implementation is very small. We then show that programming in session actors is straightforward by implementing and rewriting usecases from [4]. To our best knowledge, no other work links FSMs, actors and choreographies in a single framework. As a future work, we plan to extend to other main stream actor-based languages such as Scala and Erlang to test the generality of our framework. Also a number of additional research items regarding the expressiveness of session types still need to be carried out. Among them: how round-robin communication (as found in the rest of the usecases in [21]) can be addressed, and ... As actor languages and frameworks are becoming more popular, we believe that our work would offer an important step for writing correct large-scale actor-based communication programs.

**Acknowledgement.** We thank Ask Solem for his support and guidance. We also thank the reviewers for their valuable comments on the early version of this article. This work has been partially sponsored by VMWare, Pivotal and EPSRC EP/K011715/1, EP/K034413/1 and EP/L00058X/1, and EU project FP7-612985 (UpScale).

## REFERENCES

- [1] Cell - actors for celery. <http://cell.readthedocs.org/>.
- [2] G. Agha. *Actors: a model of concurrent computation in distributed systems*. MIT Press, Cambridge, MA, USA, 1986.
- [3] G. Agha, I. A. Mason, S. F. Smith, and C. L. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7:1–72, 1 1997.
- [4] Akka - scala actor library. <http://akka.io/>.
- [5] Advanced Message Queuing Protocol homepage. <http://www.amqp.org/>.
- [6] L. Bettini, M. Coppo, L. D’Antoni, M. D. Luca, M. Dezani-Ciancaglini, and N. Yoshida. Global progress in dynamically interleaved multiparty sessions. In *CONCUR*, volume 5201 of *LNCS*, pages 418–433. Springer, 2008.
- [7] L. Bocchi, K. Honda, E. Tuosto, and N. Yoshida. A theory of design-by-contract for distributed multiparty interactions. In *CONCUR*, volume 6269 of *LNCS*, pages 162–176, 2010.
- [8] L. Caires. Spatial-behavioral types for concurrency and resource control in distributed systems. *Theor. Comput. Sci.*, 402(2-3):120–141, 2008.
- [9] Celery. <http://http://www.celeryproject.org/>.
- [10] S. Crafa. Behavioural types for actor systems. [arXiv:1206.1687](https://arxiv.org/abs/1206.1687).
- [11] R. Demangeon and K. Honda. Nested protocols in session types. In *CONCUR*, volume 7454 of *LNCS*, pages 272–286. Springer, 2012.

- [12] R. Demangeon, K. Honda, R. Hu, R. Neykova, and N. Yoshida. Practical interruptible conversations: Distributed dynamic verification with multiparty session types and python. *FMSD*, pages 1–29, 2015.
- [13] P.-M. Deniérou and N. Yoshida. Dynamic multirole session types. In *POPL*, pages 435–446, 2011.
- [14] C. Fournet and G. Gonthier. The join calculus: A language for distributed mobile programming. In *Applied Semantics, International Summer School, APPSEM, LNCS*, pages 268–332. Springer, 2000.
- [15] S. J. Gay, V. T. Vasconcelos, A. Ravara, N. Gesbert, and A. Z. Caldeira. Modular session types for distributed object-oriented programming. In *POPL*, pages 299–312. ACM, 2010.
- [16] L. Henrio, F. Huet, and Z. István. Multi-threaded active objects. In *COORDINATION*, volume 7890 of *LNCS*, pages 90–104. Springer, 2013.
- [17] K. Honda, R. Hu, R. Neykova, T.-C. Chen, R. Demangeon, P.-M. Denilou, and N. Yoshida. Structuring communication with session types. In *COB'14*, volume 8665 of *LNCS*, pages 105–127. Springer, 2014.
- [18] K. Honda, A. Mukhamedov, G. Brown, T.-C. Chen, and N. Yoshida. Scribbling interactions with a formal foundation. In *ICDCIT*, volume 6536 of *LNCS*, pages 55–75. Springer, 2011.
- [19] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *POPL'08*, pages 273–284. ACM, 2008.
- [20] R. Hu, R. Neykova, N. Yoshida, and R. Demangeon. Practical Interruptible Conversations: Distributed Dynamic Verification with Session Types and Python. In *RV'13*, volume 8174 of *LNCS*, pages 130–148. Springer, 2013.
- [21] S. Imam and V. Sarkar. Savina - An Actor Benchmark Suite. In *4th International Workshop on Programming based on Actors, Agents, and Decentralized Control*, AGERE! 2014, October 2014.
- [22] S. M. Imam and V. Sarkar. Integrating task parallelism with actors. *SIGPLAN Not.*, 47(10):753–772, Oct. 2012.
- [23] D. Mostrous and V. T. Vasconcelos. Session Typing for a Featherweight Erlang. In *COORDINATION*, volume 6721 of *LNCS*, pages 95–109. Springer, 2011.
- [24] R. Neykova and N. Yoshida. Multiparty session actors. In eva Kühn and R. Pugliese, editors, *COORDINATION 2014*, volume 8459 of *LNCS*, pages 131–146. Springer, 2014.
- [25] R. Neykova, N. Yoshida, and R. Hu. SPY: Local Verification of Global Protocols. In *RV'13*, volume 8174 of *LNCS*, pages 358–363. Springer, 2013.
- [26] N. Ng and N. Yoshida. Pabble: Parameterised scribble for parallel programming. In *22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP*, pages 707–714. IEEE, 2014.
- [27] Ocean Observatories Initiative. <http://www.oceanobservatories.org/>.
- [28] S. Ren, Y. Yu, N. Chen, K. Marth, P.-E. Poirot, and L. Shen. Actors, roles and coordinators - a coordination model for open distributed and embedded systems. In *COORDINATION*, volume 4038 of *LNCS*, pages 247–265. Springer, 2006.
- [29] Savina source. <https://github.com/shamsmahmood/savina>.
- [30] J. Schäfer and A. Poetzsch-Heffter. Jacobox: Generalizing active objects to concurrent components. In *ECOOP*, pages 275–299, Berlin, Heidelberg, 2010. Springer-Verlag.
- [31] C. Scholliers, É. Tanter, and W. D. Meuter. Parallel actor monitors: Disentangling task-level parallelism from data partitioning in the actor model. *Sci. Comput. Program.*, 80:52–64, 2014.
- [32] Scribble project home page. <http://www.scribble.org>.
- [33] Online appendix for this article. <http://www.doc.ic.ac.uk/~rn710/sactor>.
- [34] M. Sulzmann, E. S. L. Lam, and P. Van Weert. Actors with multi-headed message receive patterns. In *COORDINATION*, pages 315–330, Berlin, Heidelberg, 2008. Springer-Verlag.
- [35] S. Tasharofi, P. Dinges, and R. E. Johnson. Why do Scala developers mix the actor model with other concurrency models? In *ECOOP*, volume 7920, pages 302–326. Springer, 2013.
- [36] C. Varela and G. Agha. Programming dynamically reconfigurable open systems with salsa. *SIGPLAN Not.*, 36(12):20–34, Dec. 2001.
- [37] BETTY WG3 - Languages Survey. [http://www.doc.ic.ac.uk/~yoshida/WG3/BETTY\\_WG3\\_state\\_of\\_art.pdf](http://www.doc.ic.ac.uk/~yoshida/WG3/BETTY_WG3_state_of_art.pdf).

#### APPENDIX A. IMPLEMENTATIONS

Here we present the corresponding Python code for some of the example.

We give in Listing 14 the implementation of the Fibonacci protocol. The `first` argument in the initialisation of a `Fibonacci` actor specifies if this is the initial actor to receive the number to be computed. As prescribed by the protocol, two message handlers, `request` and `response`, are implemented. The latter is the message received by the parent role `P`, while the former is received from the child role `Ch`. Line 13–17 illustrate a new method on the `ActorRole` class for

LISTING 14. Fibonacci in Session Actors

```

1  @protocol(c, Fibonacci, P, Ch)
2  class Fibonacci(Actor):
3      def __init__(self, first = False):
4          self.ready = false
5          self.current = 0
6          self.first = first
7
8      @role(c, P)
9      def request(self, c, n):
10         if n<=2:
11             c.P.send.response(1)
12         else:
13             fib1 = c.subprotocol("Fibonacci",
14                 P=self, Ch=new Fibonacci())
15             fib1.C.send.request(n-1)
16
17             fib2 = c.subprotocol("Fibonacci",
18                 P=self, Ch=new Fibonacci())
19             fib2.C.send.request(n-2)
20
21     @role(c, Ch)
22     def response(self, c, n):
23         if self.ready:
24             if self.is_first:
25                 print 'The result is', n + self.current
26             else:
27                 c.P.send.response(n + self.current)
28         else:
29             self.ready = True
30             self.current = n

```

creating subprotocols. Note that the `response` method should be called twice (once for each subprotocol). When the first response message is received the payload is saved temporary in the actor field `self.current` (line 30). When the second response is received, its value is added to `self.current` and the sum is returned to the parent actor (line 27).

**Counter.** We give in Listing 15 the implementation of the Counter protocol. We have two Actor classes, `Counter` and `Producer`. The interactions start with the `join` method of the `Producer` class. The body defines consecutive sending of  $n$  messages (invoking the message handler `val`) of the `Counter` actor (role `C` in the protocol `CounterProtocol`). In the `val` method the counter actor accumulates the number of received messages as part of its internal state (`self.current`). After the iteration is completed, the `Producer` sends `retrieve_message` to the `Counter`, asking for the result (line 20). The annotation on `retrieve_message` specifies that the receiver of the message is the role instance `cc` and the sender is the role `P` (`Producer`). In the body of `retrieve_message`, the `Counter` sends to `P` (`cc.P.res`) the final result (`self.current`).

**DiningPhilosopher.** We give in Listing 16 the implementation of the DiningPhilosophers protocol. For sake of space and clarity, the implementation shows only the interactions for the `Arbitrer` and `Philospher`, abstracting from the resource management done by the `Arbitrer` (in `free_forks()`). For a `Philospher` the protocol starts from the method `join`. Using its role instance, the actor does the first request to the `Arbitrer` to eat (`c.A.req(self.no)`). To distinguish between `Philosphers`, each one is given a number when initialised and when a `Philospher` sends a message to an `Arbitrer` (`c.A.req`) it also sends its number. The `Arbitrer` uses this information to identify `Philospher` that is waiting for the reply (`c.Ph[ph_no]`). Note that in line 26 and in line 28 the `Arbitrer` does

## LISTING 15. Counter in Session Actors

```

1  @protocol(cc, Counter, C, P)
2  class Counter(Actor):
3      def __init__(self):
4          self.current = 0
5
6      @role(cc)
7      def val(self, n):
8          self.current+=n
9
10     @role(cc, P)
11     def retrieve_message(self, cc):
12         cc.P.send.res(self.current)
13
14 @protocol(cc, CounterProtocol, C, P)
15 class Producer(Actor):
16     @role(cc, self)
17     def join(n, cc):
18         for i in range(1, n):
19             cc.C.send.val(1)
20             cc.C.send.retrieve_message()
21
22     @role(cc, C)
23     def res(self, n)
24         print 'The result is', n

```

not know which Philosopher is to receive a message (the annotation only specifies that the sender is a role of type Ph (Philosopher)). When a Philosopher receives a yes reply it sends a message to itself to eat (`self.become(self.eat)`). If no is returned (line 17), the Philosopher repeats its request to the Arbitrer (`c.A.req`).

**CigaretteSmoker.** We give in Listing 17 the implementation of the CigaretteSmoker protocol. Similar to the DiningPhilosopher, the protocol is initialised by giving the number of participating of Smokers, e.g  $n$ . The Arbitrer initiates the first interaction through its `join` method, where it sends a message `start_smoke` to a random smoker (it chooses random smoker from the list of smokers—`S[random]`). The Smoker actor has a handler for `start_smoke`. The handler is annotated with a `@role`, specifying this is a handler for the role smoker `c` and the caller for that handler will be the role `A`. Then the Smoker notifies the Arbitrer so the Arbitrer can proceed and give resources to other Smokers. When the Arbitrer decides to end the smoking it sends `exit` message to everyone.

## LISTING 16. DiningPhilosphers in Session Actors

```
1 @protocol(c, DiningPhilosophers, Ph, A)
2 class Philosopher(Actor):
3     def __init__(self, no):
4         self.no = no
5
6     def join(self, c):
7         c.A.send.req(self.no)
8
9     def yes(self):
10        self.become(self.eat)
11
12    @role(c, self)
13    def eat(self):
14        if full:
15            c.A.send.done()
16
17    def no(self):
18        c.A.send.req(self.no)
19
20 @protocol(c, DiningPhilosophers, A, Ph)
21 class Arbitrator(Actor):
22
23    @role(c, Ph)
24    def req(self, c, sender=ph_no)
25        if self.free_forks():
26            c.Ph[ph_no].send.yes()
27        else:
28            c.Ph[ph_no].send.no()
29    @role(c, Ph)
30    def done(self, c):
31        print 'The dinner is over'
```

## LISTING 17. CigaretteSmoker in Session Actors

```

1  @protocol(c, CigaretteSmoker, A, S)
2  class Arbiter(Actor):
3      def __init__(self, n):
4          self.n = n
5
6      def join(self, c):
7          # sends a message to a random smokers
8          radom = generate_random()
9          c.S[random].send.start_smoking()
10
11     @role(c, S)
12     def started_smoking(self):
13         if enough:
14             for i in range(1, n):
15                 c.S[i].send.exit()
16         else:
17             radom = generate_random()
18             c.S[random].send.start_smoking()
19
20     @protocol(c, CigaretteSmoker, S, A)
21     class Smoker(Actor):
22         @role(c, A)
23         def start_smoking(self):
24             c.A.send.started_smoking()
25
26         @role(c, A)
27         def exit(self):
28             print 'I am done.'
```