

Locally Static, Globally Dynamic Session Types for Active Objects

Reiner Hähnle 

Technical University Darmstadt, Germany
reiner.haehnle@tu-darmstadt.de

Anton W. Haubner

Technical University Darmstadt, Germany
anton.haubner@stud.tu-darmstadt.de

Eduard Kamburjan 

Technical University Darmstadt, Germany
kamburjan@cs.tu-darmstadt.de

Abstract

Active object languages offer an attractive trade-off between low-level, preemptive concurrency and fully distributed actors: syntactically identifiable atomic code segments and asynchronous calls are the basis of cooperative concurrency, still permitting interleaving, but nevertheless being mechanically analyzable. The challenge is to reconcile local static analysis of atomic segments with the global scheduling constraints it depends on. Here, we propose an approximate, hybrid approach; At compile-time we perform a local static analysis: later, any run not complying to a global specification is excluded via runtime checks. That specification is expressed in a type-theoretic language inspired by session types. The approach reverses the usual (first global, then local) order of analysis and, thereby, supports analysis of open distributed systems.

2012 ACM Subject Classification Theory of computation → Distributed computing models; Theory of computation → Object oriented constructs; Theory of computation → Type structures

Keywords and phrases Session Types, Active Objects, Runtime Verification, Static Verification

Digital Object Identifier 10.4230/OASICS.Gabbrielli.2020.1

Funding This research is supported by the *Constraint-Based Operational Consistency of Evolving Software Systems* (C0CoS) project, funded by the DFG as project 351097374.

For Maurizio Gabbrielli:
“*Les raisins, ou la mort!*”

1 Introduction

Lately, programming languages based on actors and *active objects* (AO) attracted a lot of interest in both academia and industry. Active objects [11] are an object-oriented modeling formalism, extending the actor model of distributed systems [21]. One prominent representative, the *abstract behavioral specification* (ABS) [24] language, was successfully applied in a variety of domains, ranging from railway operations [31] to cloud-based systems [40].

One of the advantages of ABS is its rich analysis framework with tools based on dataflow and graph analyses [3], deductive verification [13], and behavioral types [18]. However, for the time being, there is no support for code generation from scheduling policies (except user-defined schedulers at the object level), or for runtime verification beyond simple **assert** statements. The reasons lie in the AO (ABS) concurrency model.

Communication between Active Objects. An Active Object is a strongly encapsulated entity whose fields can only be accessed by getter and setter methods. Like an object in standard OO, an AO declares a set of methods, including constructors. Its peculiarity is that



© Reiner Hähnle, Anton W. Haubner and Eduard Kamburjan;
licensed under Creative Commons License CC-BY

Recent Developments in the Design and Implementation of Programming Languages.

Editors: Frank S. de Boer and Jacopo Mauro; Article No. 1; pp. 1:1–1:24

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

each method consists of *syntactically* marked atomic segments whose execution cannot be preempted. At most one task, executing the code of an atomic segment, is active at any time on the object's single processor. The advantage is that each atomic segment functions as a sequential program and can be analyzed (method-) *locally* in a *modular* fashion. However, its scheduling condition may depend on the availability of results provided by other methods, not necessarily from the same object. But such synchronization patterns can only be understood from an *object*-local or even *global* (program-wide) perspective. This is bad news for the local analysis of atomic segments as well, because in general they require information about previously scheduled tasks in order to guarantee meaningful properties. This dependency is an obstacle to any modular, global analysis. While it is possible to write sufficiently strong local contracts [30], it is a difficult, manual task, which does not align well with a top-down design starting from global communication patterns. In consequence, the ability to verify closed systems (that do not interact with their environment) is limited. Even worse, the analysis of open ABS models is generally impossible.

Additionally to cooperative method contracts [30], ABS currently uses Session Types [28, 29] for the verification of communication patterns. Both approaches consist of a *local* part that analyzes the code of single methods and a *global* part, feeding into it, that analyzes scheduling, synchronization, and messages. The global part is significantly more imprecise than the local one, because it abstracts away from functional behavior. The local specification is related to the global specification via a process called projection (from objects down to methods and atomic segments): the composition principle of the analysis follows the composition principle of the AO concurrency model [19].

Locally Static, Globally Dynamic Approach. In this paper we reverse the analysis sequence and partially move it from compile-time to runtime, resulting in a *hybrid* verification method. Specifically, local analysis is done statically, at compile-time, while global analysis is performed *later* at runtime. This is achieved by a modification of the workflow of session types: Classically, projection ensures that messages always arrive in their correct order. We retain projection, but only infer the *correct message order per object*, then construct a scheduler that enforces this order.

Local static checks permit to verify open systems: a locally specified ABS model provides only methods that perform locally correct steps, while at runtime it is ensured that methods are called correctly and in correct order. The downside is, obviously, that global errors are only detected at runtime, however, in an open system this is the only option. The second limitation of our approach is that it is not designed to perform full functional verification of state invariants, unlike interactive, deductive verification [30]. We aim at a lightweight, fully automatic method that nevertheless allows to express non-trivial properties and facilitates top-down design of distributed systems.

Yet, our approach does not modify the ABS concurrency model and requires as the single extension the availability of user-defined schedulers, i.e., the ability to reject certain task sequences. From the point of modularity, we can now verify *object*-local behavior. We implemented and evaluated our approach and illustrate with a case study that it is possible to ensure an open system always follows a given protocol.

Structure. In Sect. 2 we introduce active objects, ABS, and a suitable notion of session types. Sect. 3 describes scheduler generation and instrumentation, Sect. 4 describes and evaluates the implementation. In Sect. 5 we discuss related approaches. Finally, Sect. 6 concludes and gives future work. For space reasons, here we can describe the main ideas only with a limited degree of precision. A fully formal treatment is found in the report [20].

2 Active Objects and Session Types

The concurrency model of AO, as explained above, rests on a syntactically identifiable notion of atomic code segments that cannot be preempted. Together with strong encapsulation, this ensures that an object's state can only be modified by its own methods (including setters) and any state change must adhere to the local specification of an atomic segment. This is the basis to establish an object's invariant by suitable, *cooperative scheduling* of methods and their atomic segments. To make this work it is necessary to call a method of another (or even the own) object *without blocking*.

All active object languages, therefore, feature non-blocking, *asynchronous* method calls that return a future [4], a handle to the task executing the call.

$\text{Prgm} ::= \overrightarrow{\text{ID}} \overrightarrow{\text{CD}} \text{Main}$	$\text{ID} ::= \text{interface } I [\text{extends } \overrightarrow{I}]? \{\overrightarrow{\text{MS}}\}$	Program, Interfaces
$\text{CD} ::= \text{class } c [\text{implements } \overrightarrow{I}]? [(\overrightarrow{\text{T}} \text{ f})]? \{\overrightarrow{\text{FD}} \overrightarrow{\text{Met}} \text{Run}\}$	$\text{Main} ::= \{s\}$	Classes, Main
$\text{Run} ::= \text{Unit run}() \{s\}$	$\text{FD} ::= \text{T f} = e$	Run Method, Fields
$\text{MS} ::= \text{T m}(\overrightarrow{\text{T}} \overrightarrow{\text{v}})$	$\text{Met} ::= \text{MS } \{s; \text{return } e;\}$	Signatures, Methods
$s ::= \text{while } (e) \{s\} \mid \text{if } (e) \{s\} [\text{else } \{s\}]? \mid s; s$		Statements
$\mid \text{case } (e) \{\overrightarrow{e \Rightarrow s};\} \mid \text{await } g \mid [\text{T? } e]? = \text{rhs}$		
$g ::= e? \quad \text{rhs} ::= e \mid \text{new } C(\overrightarrow{e}) \mid e.\text{get} \mid e!\text{m}(\overrightarrow{e})$		Guards and RHS's

Figure 1 ABS grammar. T ranges over types, I over interfaces and C over classes.

The various AO languages differ in the details of how synchronization is performed, so we now turn to their specific realization in ABS. The syntax of ABS is given by the grammar in Fig. 1. With e we denote standard expressions over fields f , variables v and operators $!$, $\&$, $>=$, $<=$, $+$, $-$, $*$, $/$. Additionally, we use an expression **destiny** to access the currently computed future. Types T are all interface names (ABS enforces programming to interfaces), type-generic futures **Fut** $\langle T \rangle$, lists **List** $\langle T \rangle$, **Int**, **Unit**, and **Bool**. We also assume the usual functions for lists, etc.

In the final expression of the rule for **rhs**, the syntax for asynchronous method calls is shown (for simplicity, we leave out standard synchronous calls). As usual, a “!” replaces the dot. Asynchronous calls are always executable. Their result is a *future* of type **Fut** $\langle T \rangle$, where T is the return type of m . The *effect* of an asynchronous call is to create a task to execute m 's body in e 's object o , to be scheduled at some time in the future. In case o is also the caller, obviously the calling method must first suspend, before the callee can be scheduled. Asynchronous calls occur only as right-hand side expressions, so the future is stored in a field (or variable) f . Once the result of the computation performed by $m(\overrightarrow{e})$ is ready, it can be retrieved with the expression $f.\text{get}$. If the result is *not* ready, the **get** expression blocks the calling object. This can easily lead to a deadlock, so one typically *guards* a **get** expression with an **await** statement of the form **await** $e?$, where e 's type is of the form **Fut** $\langle T \rangle$. The effect is that execution of the current task is suspended and only rescheduled after the result of e is ready. The **await** statement and the syntactic end of a method block *are the only places, where task suspension in ABS can occur*. This justifies the following definition:

► **Definition 1** (Atomic Segment). *Code sequences starting either at the syntactic beginning of a method body or at the statement right after an **await** statement and ending either at the syntactic end of a method body or with an **await** statement, such that they contain at most the **await** statement at the end, are called atomic segment.*

Generally, ABS programs follow a simple, but standard OO paradigm in any other aspect: A program contains a main method `Main`, interfaces \overrightarrow{ID} and classes \overrightarrow{CD} . Interfaces are standard, the main method contains a list of object creations. Classes can have parameters \overrightarrow{Tf} , these are fields being initialized during object creation. The parameter type may have the form $\overrightarrow{Fut}\langle T \rangle$, i.e., futures may be passed as arguments to other methods. Classes have fields \overrightarrow{FD} , methods \overrightarrow{Met} , and a run method `Run` to start a process.

► **Example 2.** Let us illustrate cooperative scheduling in ABS with the example in Fig. 2. The program models the behavior of a mail server with notification service. It consists of three objects created in the main block: a mail server `m`, a user interface `u`, and a notification service `n`, which knows both the mail server and the user interface. Then the notification service's only method `init()` is run on `n`. The method has a single loop that periodically checks whether mail arrived and, if this is the case, notifies the user via interface `u`. Checking for mail and notifying the user require asynchronous calls to `m` and `u`, respectively, so we allocate suitable fields `fCheckMail` and `fPopup` of future type. Checking the mail must be finished before notification is handled. This is ensured by the `await` statement in Line 12. At this point, `init()` suspends. In the example, `init()` is the only method executing on `m`, so the processor will be simply idle, but it is conceivable that the main method starts other tasks on the mail server which at this point can be interleaved. Checking for mail is modelled by randomly choosing one of the literals `Mail` or `NoMail` as a return value.

Once the `response` is available, the user is notified in case there is mail, otherwise, nothing happens. Since the call to `popup()` is asynchronous, in the absence of a defined scheduler, the sequence of multiple calls to `popup()` is not necessarily in the order of mail arriving. However, the code ensures that the number of completed or active calls to `popup()` is always less than the completed calls to `checkMail()`, that there can be at most one call to `popup()` between any two calls to `checkMail()`, etc. We will show that session types are suitable to specify such global behavior in a succinct way, which then can be enforced at runtime.

Before we define session types for AO, we need to set up the machinery of user-defined schedulers needed to implement runtime checks.

A user-defined scheduler [5] is a side-effect free function in ABS that takes as parameters (1) a list of schedulable *processes* and (2) several fields of its class. It returns either `Nothing` or `Just(p)`, where `p` is one of the processes in the input list. The return value controls scheduling: if `Nothing` is returned, no process is scheduled, otherwise the chosen process is scheduled next. A process is represented as an abstract data type `Process`, i.e., an ADT that cannot be constructed manually. Instead one can access the future of a process with `destinyOf(p)` and its methodname as a `String` with `method(p)`.

Lists are also ADTs and `nth(input,i)` returns the *i*-th element. Keyword `def` is used to define a function with parameters that evaluates a result using standard expressions (and recursion). ABS does not support fully-fledged functional programming and only a fixed set of higher-order functions. For example, higher-order functions such as `map` and `filter` are part of the ABS standard library.

► **Example 3.** Let us consider the following scheduler and class.

```
def Maybe<Process> scheduler(List<Process> input, Int y, String m) =
  if ( y < 0 || y >= length(input) ) Nothing else
    if ( method(nth(input,y)) != m ) Just(nth(input,y)) else Nothing;

[Scheduler: scheduler(queue, y, m)]
class C(String m, Int y) { ... }
```

```

1 data Msg = Mail | NoMail;
2
3 class NotifyService
4   (MailServerI m, UII u)
5   implements NotifyServiceI {
6   Fut<Msg> fCheckMail;
7   Fut<Unit> fPopup;
8   Unit init(int bound) {
9     Int i = 0;
10    while (i < bound) {
11      fCheckMail = m!checkMail();
12      await fCheckMail?;
13      Msg response = fCheckMail.get;
14      case (response) {
15        Mail => fPopup = u!popup(i);
16        NoMail => skip;
17      }
18      i = i + 1;
19    }
20  }
21 }

22 class MailServer
23   implements MailServerI {
24   Msg checkMail() {
25     Msg result = NoMail;
26     if (random(2) == 1)
27       result = Mail;
28     return result;
29   }
30 }
31 class UI implements UII {
32   Unit popup(int id) {
33     println("You got mail! Id " + id);
34   }
35 }
36 { // Main block
37   MailServerI m = new MailServer();
38   UII u = new UI();
39   NotifyServiceI n
40     = new NotifyService(m,u);
41   await n!init(42);
42 }

```

■ Figure 2 Mail server example in ABS.

167 y and m, the field names and their types, must be identical in class and scheduler function
 168 to use the scheduler. The code above selects the y-th element in the input list, unless it is out
 169 of range or a process executing a method named m. The annotation [Scheduler: scheduler(
 170 queue, y, m)] connects scheduler and class. Whenever the scheduler is invoked, the input
 171 list is guaranteed to be non-empty.

172 2.1 Session Types for Active Objects

173 Session types specify and verify the behavior of a closed unit of communication, called
 174 a *session*. A session type specification consists of three parts: (1) global types, a global
 175 specification of the session, (2) local types, specifications for the endpoints in a session, and
 176 (3) a projection mechanism that generates a local specification for each endpoint participating
 177 in the communication from a global type. For checking that the whole unit adheres to its
 178 global specification, it suffices to check the local endpoints and, possibly, side-conditions on
 179 the unit. Additionally, the session type system needs some kind of mechanism to ensure that
 180 the local endpoints adhere to their local type.

181 In the original formulation for the π -calculus [8] a session is centered around a channel,
 182 endpoints are the processes participating in the communication over the typed channel.
 183 To ensure that projection succeeds, a linearity check on the channel is performed as a
 184 side-condition of projection.

185 For Active Objects, the situation changes: there are no channels and endpoints partic-
 186 ipating in any communication are not uniform, because the target of a method call is an
 187 object, but the target of a future read is a (terminated) process. As there are no channels, a
 188 linearity check to ensure that messages arrive in the right order is impossible.

189 Session types for AO [28, 29] adopt and adapt the concepts of session types for channels:
 190 **Unit of Communication:** The unit of communication is described by a set of objects that

only contain pointers to each other.

Endpoints: The notion of endpoint is two-fold. Both objects and processes are endpoints and the projection of a global session type first projects on the object and then projects one more time on the processes inside that object. The result of the first projection is an *object-local* type and the result of the second projection is a *method-local* type.¹

Order of Messages: To ensure messages arrive in the correct order, a static analysis can be used to determine whether the order of messages is total from the perspective of each object (but not globally).

Here we remove the check on message order at the level of the type system and instead enforce it at runtime using the structure provided by the object-local type. Before we introduce syntax of global and local types, it is worth mentioning that we are only concerned with *protocol adherence*: Does the system implement the protocol described by the global type? We ignore *deadlock freedom*, which can be approached either with session types [29] or a dedicated deadlock checker for AO [15, 18, 25]. The system we introducing below is a slight variation of the session types in [29].

2.1.1 Global Types

Global types follow the structure of regular expressions and allow Kleene star-style repetition, sequence and branching. Branching is guarded by a single role that determines which branch of the protocol to follow. As single actions, the type defines a certain kind of interaction between two roles or a role and a process/future. To keep track of processes and futures within a protocol, we use tracked futures: references to the future of a specified method call.

► **Definition 4.** Let \mathbf{p}, \mathbf{q} range over roles, \mathbf{t} over tracked futures and \mathbf{C} over ADT constructors. The syntax of global protocols \mathbf{GP} and global types \mathbf{G} is defined in Fig. 3. Specifications $\langle \cdot \rangle$ are all optional.

$\mathbf{GP} ::= \mathbf{0} \xrightarrow{\mathbf{t}} \mathbf{p} : \mathbf{m} . \mathbf{G}$	Global Protocol
$\mathbf{G} ::= \mathbf{p} \xrightarrow{\mathbf{t}} \mathbf{q} : \mathbf{m} \mid \mathbf{p} \downarrow \mathbf{t}(\mathbf{C}) \mid \mathbf{p} \uparrow \mathbf{t}(\mathbf{C})$	Call, Termination and Synchronization Action
$\mid \text{Rel}(\mathbf{p}, \mathbf{t}) \mid \text{skip}$	Suspension and Empty Action
$\mid \mathbf{p}\{\mathbf{G}_i\}_{i \in I} \mid (\mathbf{G})^* \mid \mathbf{G} . \mathbf{G}$	Branching, Repetition and Sequential Composition

■ **Figure 3** Syntax of Global Session Types

The global protocol starts with $\mathbf{0} \xrightarrow{\mathbf{t}} \mathbf{p} : \mathbf{m}$ and specifies how the session is started. The call action $\mathbf{p} \xrightarrow{\mathbf{t}} \mathbf{q} : \mathbf{m}$ specifies a call from the object with role \mathbf{p} to the one with role \mathbf{q} on method \mathbf{m} . This process is tracked by \mathbf{t} in the rest of the type. The termination action $\mathbf{p} \downarrow \mathbf{t}(\mathbf{C})$ specifies that the object with role \mathbf{p} terminates the process tracked by \mathbf{t} and the return value has the outermost constructor \mathbf{C} . The synchronization action $\mathbf{p} \uparrow \mathbf{t}(\mathbf{C})$ specifies that the object with role \mathbf{p} reads from the future tracked by \mathbf{t} and reads a value with the outermost constructor \mathbf{C} . The suspension action $\text{Rel}(\mathbf{p}, \mathbf{t})$ specifies that the object with role \mathbf{p} suspends its currently active process until the future tracked by \mathbf{t} is resolved. We stress

¹ The projection may also be done in one step [27], but this removes the object-local types which we are investigating in this paper.

that t is *not* the tracked future of the suspended process. The empty action specifies no action and is needed to specify, e.g., branches without visible actions. Branching $\mathbf{p}\{\mathbf{G}_i\}_{i \in I}$ specifies that the object with role \mathbf{p} chooses one of the branches \mathbf{G}_i to continue the protocol. Finally, repetition and sequential composition are analogous to regular expressions.

► **Example 5.** We continue Ex. 2. The roles of the protocol are named **NS** for the notification service, **Mail** for the mail server and **GUI** for the GUI. The intended behavior is specified by the following global type:

$$\begin{aligned} & \mathbf{0} \xrightarrow{t_0} \mathbf{NS} : \text{init} . \\ & \left(\mathbf{NS} \xrightarrow{t_1} \mathbf{Mail} : \text{check} . \text{Rel}(\mathbf{NS}, t_1) . \right. \\ & \quad \mathbf{Mail} \left\{ \begin{array}{l} \mathbf{Mail} \downarrow t_1 \langle \text{NewMail} \rangle . \mathbf{NS} \uparrow t_1 \langle \text{NewMail} \rangle . \mathbf{NS} \xrightarrow{t_2} \mathbf{GUI} : \text{show} . \mathbf{GUI} \downarrow t_2 \\ \mathbf{Mail} \downarrow t_1 \langle \text{NoMail} \rangle . \mathbf{NS} \uparrow t_1 \langle \text{NoMail} \rangle \end{array} \right\} \\ & \left. \right)^* . \mathbf{NS} \downarrow t_0 \end{aligned}$$

The above example demonstrates the use of tracked futures and repetition, but it is strongly synchronized: The described synchronization structure enforces correct interaction order, no deviation due to the scheduler is possible. In contrast, consider the following scenario and global session type that is not strongly synchronized.

► **Example 6.** The protocol describes four roles: a student **S**, a service desk **D**, a computation server **C** and a report generator **R**. The computation server computes the grade of a student, and sends it to the report generator, which in turn generates a report that is send to the service desk. The computation server notifies the student that its grade has been computed. The service desk may only serve the student after the report has arrived. This is specified by the following global protocol:

$$\begin{aligned} & \mathbf{0} \xrightarrow{t_0} \mathbf{C} : \text{compute} . \mathbf{C} \xrightarrow{t_1} \mathbf{R} : \text{toReport} . \mathbf{C} \xrightarrow{t_2} \mathbf{S} : \text{notify} . \mathbf{R} \xrightarrow{t_3} \mathbf{D} : \text{publish} . \mathbf{R} \downarrow t_1 \\ & . \mathbf{D} \downarrow t_3 . \mathbf{S} \xrightarrow{t_4} \mathbf{D} : \text{request} . \mathbf{D} \downarrow t_4 . \mathbf{S} \uparrow t_4 . \mathbf{S} \downarrow t_2 . \mathbf{C} \downarrow t_0 \end{aligned}$$

Note that **D** is called on **request** and **publish** but no synchronization ensures that those messages arrive in the specified order.

2.1.2 Local Types

We distinguish object-local types, method-local types and scheduling types. Method-local and object-local types differ syntactically only in their passive choice operator and the specification of synchronization. Scheduling types describe the actions of the scheduler of a role and share their syntax with method-local types.

► **Definition 7.** Let \mathbf{p} range over roles and $\mathbf{0}, \mathbf{m}$ over method names, \mathbf{t} over tracked futures and \mathbf{C} over ADT constructors. The syntax of object-local types \mathbf{L} is defined as follows:

$\mathbf{L} ::= \mathbf{p} ?_{\mathbf{t}} \mathbf{m} \mid \mathbf{p} !_{\mathbf{t}} \mathbf{m} \mid \text{Put } \mathbf{t}(\mathbf{C})$	Receiving, Sending and Termination Action
$\text{Get } \mathbf{t}(\mathbf{C}) \mid \text{Susp}(\mathbf{t}, \mathbf{t}) \mid \text{React } \mathbf{t}$	Synchronization, Suspension and Reactivation Action
$\&_{\mathbf{t}} \{\mathbf{L}_i\}_{i \in I} \mid \oplus \{\mathbf{L}_i\}_{i \in I}$	Passive and Active Choice
$\text{skip} \mid \mathbf{L} . \mathbf{L} \mid (\mathbf{L})^*$	Empty Action, Sequential Composition and Repetition

The syntax of method-local types is analogous, but (1) the synchronization action takes no \mathbf{C} specification and (2) passive choice takes the following form, called guarded passive choice:

$$\&_{\mathbf{t}} \{\mathbf{C}_i : \mathbf{L}_i\}_{i \in I}$$

The receiving action is the callee's view on the global call action, \mathbf{p} is the caller. Note that a method in **ABS** has no access to the caller, but we may access it in the scheduler. The sending action is the caller's view on the global call action, \mathbf{p} is the callee. The local termination action is the equivalent of the global termination action. The local suspension action $\text{Susp}(t_1, t_2)$ specifies that the process computing t_1 suspends until t_2 is known. The reactivation action $\text{React } t$ specifies reactivation of the process computing t . These two actions are the local view on the global suspension action, but (1) locally the suspending process is known and (2) one can infer, where the reactivation must happen. We use three choice operators:

- (Object-Local) Unguarded passive choice $\&_t\{\mathbf{L}_i\}_{i \in I}$ specifies that the object reacts to the choice stored in t . The choice is stored as the \mathbf{C} parameter of the first **Get** action on t in the given branch.
- (Method-Local) Guarded passive choice $\&_t\{\mathbf{C}_i : \mathbf{L}_i\}_{i \in I}$ specifies which constructor corresponds to which branch directly, as it is only indirectly encoded in the unguarded passive choice.
- Active choice $\oplus\{\mathbf{L}_i\}_{i \in I}$ specifies that the object or process in question chooses one of the branches to continue. It is not specified how the choice is made.²

The remaining actions are analogous to their global counterpart.

We introduce projection formally in the subsequent section, but provide examples of local types based on Ex. 5, 6 now to illustrate the differences among the various local types.

► **Example 8.** Below is the object-local type of **Mail** in Ex. 5 followed by the method-local type of t_1 and the scheduling type. The differences between the former are that (1) the method-local type contains *no* repetition, because the repetition is not visible to a single process and (2) the receiving action is omitted, because it is redundant when the tracked future is known.

$$\begin{array}{ll}
 \left(\text{NS?}_{t_1} \text{check} . \oplus \left\{ \begin{array}{l} \text{Put } t_1(\text{NewMail}) \\ \text{Put } t_1(\text{NoMail}) \end{array} \right\} \right)^* & \text{Object-local type} \\
 \oplus \left\{ \begin{array}{l} \text{Put } t_1(\text{NewMail}) \\ \text{Put } t_1(\text{NoMail}) \end{array} \right\} & \text{Method-local type} \\
 \left(\text{NS?}_{t_1} \text{check} \right)^* & \text{Scheduling type}
 \end{array}$$

The method-local type contains only the actions performed by the processes of a single method. A scheduling type contains only the actions needed for scheduling: empty, reactivation and receiving actions, as well as both kinds of branching, repetition and sequential compositions. The following is the scheduling type of **D** in Ex. 6: $\mathbf{R?}_{t_3} \text{publish} . \mathbf{S?}_{t_4} \text{request}$

3 LSGD Session Types

The verification workflow of our system takes a global type and generates an instrumented **ABS** program and a proof that each method is following its method-local type.

- First, we establish certain well-formedness conditions of the global type to ensure it describes a protocol that is realizable in the AO concurrency model.

² For guarded active choice we refer to Kamburjan & Chen [28].

- 303 ■ Then, the global type is projected on each participating object. This results in an
- 304 object-local type, describing the actions an object both expects and is obliged to perform.
- 305 ■ From the object-local type we generate (a) a session automaton that describes the order
- 306 of scheduling actions and (b) a method-local type for each method. Scheduling actions
- 307 include the receiving action (receiving method calls) and the rescheduling action (reacting
- 308 to a resolved future).
- 309 ■ The session automaton is translated into a user-defined scheduler, which is added to the
- 310 object together with fields and operations to keep track of the state.
- 311 ■ Each method is checked statically against its method-local type.

312 For brevity, we give a simplified account of the implemented system [20] and omit some
 313 features, e.g., allowing interactions with objects that do not participate in the session.

314 3.1 Session Automata

315 Before defining the workflow, we introduce Session Automata [7]. Session automata are a
 316 class of register automata [33]: finite automata over an infinite alphabet. General register
 317 automata allow to store read values of infinite alphabets in registers and compare the register
 318 contents by equality. Session automata have the restriction that only *fresh values* can be
 319 stored, i.e., values that have not been seen in the input word so far. This matches our model
 320 when futures are regarded as data and allows one to decide whether two session automata
 321 accept the same language. In our system, the alphabet is the set of futures and we only store
 322 futures upon receiving a method call. This guarantees their freshness upon storage.

323 ► **Definition 9.** *Let Σ be a finite set of labels, D an infinite set of data equipped with equality*
 324 *and $k \in \mathbb{N}$. A k -Register Session Automaton is a tuple (Q, q_0, Φ, F) , where Q is the set of*
 325 *states, $q_0 \in Q$ its start state, $F \subseteq Q$ the set of accepting states, and the transition relation is*
 326 *as follows:*

$$327 \quad \Phi \subseteq (Q \times Q) \cup (Q \times (\Sigma \times D) \times \mathcal{P}(\{1, \dots, k\}) \times \{1, \dots, k\} \times Q)$$

328 Runs of session automata are defined over stores and data words. A transition either (1)
 329 only changes the state, but neither changes the store nor consumes a letter, or (2) changes
 330 the state upon reading the next letter by comparing the data with a register in its store and
 331 storing the read data.

332 ► **Definition 10.** *A store $\sigma : \{1, \dots, k\} \mapsto D \cup \{\perp\}$ is a function from register identifiers*
 333 *to data or the special symbol \perp . The initial store σ_0 maps all register identifiers to \perp . A*
 334 *data word $w = (a_0, d_0), \dots, (a_n, d_n)$ is a finite sequence of pairs of labels and data. A run*
 335 *$(q_0, j_0, \sigma_0), \dots, (q_m, j_m, \sigma_j)$ of a k -register session automaton (Q, q_0, Φ, F) on a word w of*
 336 *length n is a sequence*

$$337 \quad s \in (Q \times \mathbb{N} \times \{1, \dots, k\} \mapsto D)^*$$


338 *where q_i is the current state, σ_i the current store and j_i the next letter. The sequence must*
 339 *start with $(q_0, 0, \sigma_0)$ and satisfy the following condition for each position $0 < i < m$:*

$$340 \quad (q_i, q_{i+1}) \in \Phi \wedge (j_i = j_{i+1}) \wedge (\sigma_i = \sigma_{i+1})$$

$$341 \quad \vee \left((q_i, (a_{j_i}, d_{j_i}), I, k, q_{i+1}) \in \Phi \wedge (j_i = j_i + 1) \wedge \sigma_{i+1} = \sigma_i[k \setminus d_{j_i}] \wedge \forall l \in I. \sigma_i(l) = d_{j_i} \right)$$

342 In the following we set $D = \text{Fut}$ and $\Sigma = \{\text{invREv}\} \times \text{Met} \cup \{\text{condREv}\}$.

► **Example 11.** The following 2-register session automaton models the scheduling type of **D** in Ex. 8. The two stores of the futures in registers r_i are used to model reactivation.



For brevity, we write $(q, (\text{invREv}, m), q')$ and $(i, (\text{condREv}), q')$ for transitions with the given label and say that register i is either written or read. We never write to or read from more than one register in a single transition.

3.2 Projection

Projection generates (1) a method-local type per participating method in the session and (2) a special object-local type, called *scheduling type*, for each role. The scheduling type describes the order of operations controlled by the scheduler, i.e., process start and rescheduling.

Projection consists of four steps: pre-analysis, projection on a role, projection on a tracked future, and generation of a scheduling type from an object-local type.

Pre-analysis: Reject obviously malformed types and annotate the global type with information used in later steps, for example, which future is currently being computed.

Projection on Role: Generate an object-local type that describes the view of a role on the global type.

Projection on Tracked Future: Generate a method-local type that describes the view of a process on the object-local type.

Generation of Scheduling Type: Generate the scheduling type that describes the operations performed by the scheduler of an object.

3.2.1 Pre-Analysis

Pre-analysis of a global type checks that it specifies a feasible protocol in the AO concurrency model. It generates an *annotated* global type $\mathbf{G}\langle\sigma\rangle$, where σ describes the specified state of a role before and after performing the specified action. We refrain from introducing all the formal details and only describe the checked properties of a global type.

Future Freshness: Each tracked future identifies exactly one call action. For example, the following type fails pre-analysis and is rejected, because t is not fresh in the second call.

$$0 \xrightarrow{t} p:m . p \xrightarrow{t} q:n . p \downarrow t . q \downarrow t$$

Actor Activity: A call action can only be specified when the callee is not specified as currently executing a method and a suspending action can only suspend a process when it is specified as being active. For example, the following global type contains two errors: the call of t_2 must wait until p is terminated and the suspension action of p cannot suspend any process.

$$0 \xrightarrow{t_0} p:m . p \xrightarrow{t_1} q:n . q \xrightarrow{t_2} p:o . p \downarrow t_0 . \text{Rel}(p, t_1) . q \downarrow t_1 . p \downarrow t_2$$

The following is one possible “debugged” version that passes pre-analysis:

$$0 \xrightarrow{t_0} p:m . p \xrightarrow{t_1} q:n . \text{Rel}(p, t_1) . q \xrightarrow{t_2} p:o . p \downarrow t_2 . q \downarrow t_1 . p \downarrow t_0$$

Resolution Analysis: A future can only be read if it has terminated before and is accessible to the reading role.³

Scope Analysis: Repetition introduces *scopes* into the specification, as a process is started exactly once and terminated exactly once. For example, the following type is not correctly scoped, because it allows situations where (1) n is never called, and thus t_1 cannot be terminated and (2) where n is called multiple times and it is not specified how many of those processes are terminated and in which order:

$$0 \xrightarrow{t_0} p : m . (p \xrightarrow{t_1} q : n)^* . p \downarrow t_1 . q \downarrow t_0$$

The scope analysis checks that (1) every tracked future that is started within a repetition is resolved within the same repetition; (2) every tracked future that is resolved within a repetition is started within the same repetition; (3) every tracked future that is synchronized upon within a repetition is started within the same repetition; (4) for every role the active tracked future and the set of suspended tracked future before and after the repetition are the same. (5) every tracked future that is started within a branch is resolved within the same branch; (6) every tracked future that is resolved within a branch is started within the same branch;

During pre-analysis each global type, except sequence, is annotated with an abstract state σ . An abstract state is a mapping from roles to a pair $(AState, SState)$, where $AState$ is either **Active**(t), expressing that the role is currently specified as executing the process for t or **Susp** if it is currently specified inactive. $SState$ is a set of pairs of tracked futures (t, t') , expressing that there is a suspended process for t waiting for t' . Pre-analysis ensures that there are no t_1, t_2 with $(t_1, t), (t_2, t) \in SState$ in any abstract state for any role, i.e., there are never two processes of one role waiting for the same future.⁴

3.2.2 Global Projection

The projection of global types on a role is defined in Fig. 4. Projection is a partial function $G \downarrow p$. It checks that any action is specified to happen when the role performing this action is active and has a process that can perform the communication. The result of projection is an object-local type, annotated with abstract states.

The initial action results in a receiving action for the callee and **skip** for any other role. Similarly, the projection of a call action is a receiving action for the callee and a sending action for the caller.

Projection of the termination action has three cases: (1) If projected on the terminating role, it is ensured that this role is active and can perform the action. The result is a local termination action. (2) If projected on a role waiting for the tracked future of the action, it is ensured that this role is inactive. The result is a reactivation action. (3) Projection on any other role results in **skip**. Projection fails if, for example, the terminating role is inactive.

Projection of synchronization results in a local synchronization action for the specified role and **skip** for any other role. It is checked that the specified role is active. The suspension action is analogous. Projection of **skip** is the identity, projection of branching results in an active choice for the specified role (which must be active) and a passive choice over the currently active future of the choosing role for any other role. Projection of the repetition

³ On passing data in Session Types for Active Objects, we refer to [27].

⁴ Because it is not *specified* in which order they should be reactivated. If such a specification were given, that order could be reflected in the projected object-local type.

$$\begin{aligned}
0 \xrightarrow{t} q : m \langle \sigma \rangle \upharpoonright p &= \begin{cases} 0 ?_t m \langle \sigma \rangle & \text{if } p = q \\ \text{skip} & \text{otherwise} \end{cases} \\
q \xrightarrow{t} r : m \langle \sigma \rangle \upharpoonright p &= \begin{cases} r !_t m \langle \sigma \rangle & \text{if } p = q \\ q ?_t m \langle \sigma \rangle & \text{if } p = r \\ \text{skip} & \text{otherwise} \end{cases} \\
q \downarrow t \langle C \rangle \langle \sigma \rangle \upharpoonright p &= \begin{cases} \text{Put } t \langle C \rangle \langle \sigma \rangle & \text{if } p = q \wedge \sigma(p) = (\text{Active}(t), SState) \\ \text{React } t' \langle \sigma \rangle & \text{if } p \neq q \wedge \sigma(p) = (\text{Susp}, SState) \wedge (t, t') \in SState \\ \text{skip} & \text{if } p \neq q \wedge \sigma(p) = (\text{Susp}, SState) \wedge \nexists t'. (t, t') \in SState \end{cases} \\
q \uparrow t \langle C \rangle \langle \sigma \rangle \upharpoonright p &= \begin{cases} \text{Get } t \langle C \rangle \langle \sigma \rangle & \text{if } p = q \wedge \sigma(p) = (\text{Active}(t'), SState) \\ \text{skip} & \text{otherwise} \end{cases} \\
\text{Rel}(q, t) \upharpoonright p &= \begin{cases} \text{Susp}(t', t) \langle \sigma \rangle & \text{if } p = q \wedge \sigma(p) = (\text{Active}(t'), SState) \\ \text{skip} & \text{if } p \neq q \end{cases} \\
q \{G_i\}_{i \in I} \langle \sigma \rangle \upharpoonright p &= \begin{cases} \oplus \{G_i \upharpoonright p \langle \sigma \rangle\}_{i \in I} & \text{if } p = q \wedge \sigma(p) = (\text{Active}(t), SState) \\ \&_t \{G_i \upharpoonright p \langle \sigma \rangle\}_{i \in I} & \text{if } p \neq q \wedge \sigma(q) = (\text{Active}(t), SState) \end{cases} \\
(G)^* \langle \sigma \rangle \upharpoonright p &= \begin{cases} (L)^* \langle \sigma \rangle & \text{if } G \upharpoonright p = L \neq \text{skip} \\ \text{skip} & \text{otherwise} \end{cases} \\
(G_1 \cdot G_2) \upharpoonright p &= (G_1 \upharpoonright p) \cdot (G_2 \upharpoonright p) \quad \text{skip} \upharpoonright p = \text{skip}
\end{aligned}$$

■ **Figure 4** Projection of global type on roles

repeats the projection of the inner part if it performs some action. Otherwise, the repetition is replaced with an empty action. Finally, projection of sequential composition is sequential composition of the projected types. We assume that structural congruence is used to remove superfluous empty actions and branching.

3.2.3 Local Projection

Local projection generates a method-local type from an object-local type for each tracked future. Each tracked future is introduced by a call action, so we can easily connect method-local types to methods. For simplicity, we demand that each method has only one type.

Local projection must invert the relation between passive choice and synchronization. A global type specifies first the choice and marks the future of the choosing role during global projection. Afterwards, the future is resolved and may be synchronized upon. Locally, however, the method synchronizes *first* and *then* branches depending on the read value. Local projection handles this by pulling out the prefix of all branches from a passive choice up to the synchronization action over the choosing future.

► **Definition 12.** Let t be a tracked future and L_i a set of object-local types. The prefix for t of some object-local L is defined as the shortest type L_t that ends in $\text{Get } t \langle C \rangle$: The function $\text{split}_t(L)$ returns the prefix and the remaining postfix of a type.

$\text{split}_t(L) = (L_{\text{head}}, C, L_{\text{tail}})$ such that

$$L_{\text{head}} = \widehat{L} \cdot \text{Get } t, \widehat{L} \text{ contains no } \text{Get } t, \text{ and } L \equiv \widehat{L} \cdot \text{Get } t \langle C \rangle \cdot L_{\text{tail}}$$

The function $\text{split}_t(\{L_i\}_{i \in I})$ returns the common prefix and the remaining postfixes of all input types. Note that the function may be undefined.

$$\text{split}_t(\{L_i\}_{i \in I}) = (L_{\text{head}}, \{(C_i, L_{\text{tail}}^i)\}) \text{ such that } \text{split}_t(L_i) = (L_{\text{head}}, C_i, L_{\text{tail}}^i)$$

Projection $L \langle \sigma \rangle \upharpoonright p$ of an annotated local type $L \langle \sigma \rangle$ on t for role p is given in Fig. 5. It removes receiving and reactivation actions, is the identity on any other non-composed action

$$\begin{aligned}
\mathbf{L}\langle\sigma\rangle \upharpoonright^{\mathbf{P}} \mathbf{t} &= \mathbf{L} \text{ if } \sigma(\mathbf{p}) = ((\text{Active}(\mathbf{t}), SState) \\
&\quad \text{and } \mathbf{L} \in \{\mathbf{p}!_{\mathbf{t}'\mathbf{m}}, \text{Put } \mathbf{t}'(\mathbf{C}), \text{Get } \mathbf{t}'(\mathbf{C}), \text{Susp}(\mathbf{t}', \mathbf{t}''), \text{skip}\} \\
\mathbf{p}?_{\mathbf{t}'\mathbf{m}} \upharpoonright^{\mathbf{P}} \mathbf{t} &= \text{React } \mathbf{t}' \upharpoonright^{\mathbf{P}} \mathbf{t} = \text{skip} \\
(\mathbf{L}_1 . \mathbf{L}_2) \upharpoonright^{\mathbf{P}} \mathbf{t} &= (\mathbf{L}_1) \upharpoonright^{\mathbf{P}} \mathbf{t} . (\mathbf{L}_2) \upharpoonright^{\mathbf{P}} \mathbf{t} \\
(\mathbf{L})^* \upharpoonright^{\mathbf{P}} \mathbf{t} &= \begin{cases} \mathbf{L} \upharpoonright^{\mathbf{P}} \mathbf{t} & \text{if } \mathbf{t} \text{ is introduced within } \mathbf{L} \\ (\mathbf{L}')^* & \text{if } \mathbf{L} \upharpoonright^{\mathbf{P}} \mathbf{t} = \mathbf{L}' \neq \text{skip} \text{ and } \mathbf{t} \text{ is not introduced within } \mathbf{L} \\ \text{skip} & \text{otherwise} \end{cases} \\
\oplus\{\mathbf{L}_i\}_{i \in I} \upharpoonright^{\mathbf{P}} \mathbf{t} &= \oplus\{\mathbf{L}_i \upharpoonright^{\mathbf{P}} \mathbf{t}\}_{i \in I} \\
&\&_{\mathbf{t}'}\{\mathbf{L}_i\}_{i \in I} \upharpoonright^{\mathbf{P}} \mathbf{t} = \begin{cases} \mathbf{L} \upharpoonright^{\mathbf{P}} \mathbf{t} . \&_{\mathbf{t}'}\{\mathbf{C}_i : \widehat{\mathbf{L}}_i \upharpoonright^{\mathbf{P}} \mathbf{t}\}_{i \in I} & \text{if } \text{split}_{\mathbf{t}'}(\{\mathbf{L}_i\}_{i \in I}) = (\mathbf{L}, \{(\mathbf{C}_i, \widehat{\mathbf{L}}_i)\}) \\ \&_{\mathbf{t}'}\{\mathbf{L}_i\}_{i \in I} \upharpoonright^{\mathbf{P}} \mathbf{t} = \mathbf{L}_j \upharpoonright^{\mathbf{P}} \mathbf{t} & \text{if } j \in J \text{ and } \mathbf{t} \text{ is introduced in } \mathbf{L}_j \end{cases}
\end{aligned}$$

■ **Figure 5** Projection of local types on tracked futures

$$\begin{aligned}
&\oplus\{\mathbf{L}_i\}_{i \in I} < \oplus\{\mathbf{L}_i\}_{i \in I \cup J} \quad \&_{\mathbf{t}'}\{\mathbf{C}_i : \mathbf{L}_i\}_{i \in I} > \&_{\mathbf{t}'}\{\mathbf{C}_i : \mathbf{L}_i\}_{i \in I \cup J} \\
&(\mathbf{L})^* < (\widehat{\mathbf{L}})^* & \text{if } \mathbf{L} < \widehat{\mathbf{L}} \\
&\mathbf{L}_1 . \mathbf{L}_2 < \widehat{\mathbf{L}}_1 . \widehat{\mathbf{L}}_2 & \text{if } \mathbf{L}_i < \widehat{\mathbf{L}}_i \\
&\oplus\{\mathbf{L}_i\}_{i \in I} < \oplus\{\widehat{\mathbf{L}}_i\}_{i \in I} & \text{if } \mathbf{L}_i < \widehat{\mathbf{L}}_i \\
&\&_{\mathbf{t}'}\{\mathbf{C}_i : \mathbf{L}_i\}_{i \in I} < \&_{\mathbf{t}'}\{\mathbf{C}_i : \widehat{\mathbf{L}}_i\}_{i \in I} & \text{if } \mathbf{L}_i < \widehat{\mathbf{L}}_i \\
&\mathbf{L} \equiv \oplus\{\mathbf{L}\} \quad \mathbf{L} \equiv \text{skip} . \mathbf{L} \quad \mathbf{L} \equiv \mathbf{L} . \text{skip} \\
&\mathbf{L} . \oplus\{\mathbf{L}_i\}_{i \in I} \equiv \oplus\{\mathbf{L} . \mathbf{L}_i\}_{i \in I} \quad \oplus\{\mathbf{L}, \text{skip}_i\}_{i \in I} \equiv \mathbf{L}
\end{aligned}$$

■ **Figure 6** Subtype relation and structural congruence of method-local types

446 and propagates on sequential composition and active choice. For passive choice, the above
 447 split is applied, unless the projection future is introduced in only one branch. Repetitions
 448 outside a single method run are removed.

449 3.2.4 Scheduling Type

450 Given a projected object-local type \mathbf{L} , the scheduling type $\mathcal{S}(\mathbf{L})$ is generated by replacing all
 451 termination, synchronization, suspension and sending actions with **skip** and using structural
 452 congruence (see Fig. 6) to simplify the result.

453 3.3 Locally Static

454 Method-local types are checked statically. This ensures that *if* every process is scheduled
 455 correctly, then the process will perform its local view on the protocol correctly. Before we
 456 present the type system itself, we define typing contexts and auxiliary functions.

457 The subtype relations $<$, \leq and structural congruence are standard, see Fig. 6. Structural
 458 congruence allows to add and remove **skip** actions. An active choice with a single branch can
 459 be simplified to the content of the branch. The interesting rules for subtyping are the ones
 460 for branching: Active branching may drop branches, as the implementing role may never
 461 take a subset of its possible choices. Its dual, passive branching, may add branches instead.

462 We use two typing contexts: Δ maps locations (fields and variables) to roles, Γ maps

463 tracked futures to pairs of locations or the symbol \perp . the Δ context ensures that a method
 464 interacts with the correct endpoints, while Γ keeps track of futures and their read values.
 465 We use some auxiliary functions and predicates:

- 466 ■ The function $\Gamma^{\mathfrak{A}}$ removes all fields from the pairs in the image of Γ .
- 467 ■ The function $\text{constr}(\mathbf{e})$ returns the outermost constructor of expression \mathbf{e} .
- 468 ■ The function $\text{def}(\mathbf{c})$ returns the declaration of class \mathbf{c} .
- 469 ■ The predicate $\text{inter}(\mathbf{s}, \Gamma)$ holds if the statement \mathbf{s} contains no **get**, no **return**, no **await**, and
 470 writes into no location that is in a pair in the image of Γ .
- 471 ■ The predicate $\mathbf{p} \in \mathbf{G}$ or $\mathbf{p} \in \mathbf{L}$ holds if the role \mathbf{p} occurs in the type.
- 472 ■ The predicate $\mathbf{e} \in \text{im}\Gamma$ holds if the location \mathbf{e} is in any pair in the image of Γ .

473 The type system is shown in Fig. 7. Rule T-main checks that the main block sets up the
 474 session correctly: Each role is assigned to exactly one object and the corresponding class is
 475 checked against the projected type on this role. Also, each parameter of a class is assigned
 476 such that the passed variable has the correct role (the f_{ij} are the fields declared in $\text{def}(\mathbf{c}_i)$).
 477 Lastly, the sole called method is correctly specified and called on the correct object. The
 478 rule T-class checks that each role needed for the object-local type is available in some field
 479 and checks each method against its method-local type.

$$\begin{array}{c}
 \text{T-main} \frac{\forall i. \exists \mathbf{p} \in \mathbf{G}. \Delta(\mathbf{v}_i) = \mathbf{p} \quad \forall \mathbf{p} \in \mathbf{G}. \exists i. \Delta(\mathbf{v}_i) = \mathbf{p} \quad \Delta_i(f_{ij}) = \Delta(\mathbf{v}_{ij}) \quad \Delta_i \vdash \text{def}(\mathbf{c}_i) : \mathbf{G} \upharpoonright \Delta(\mathbf{v}_i) \quad \Delta(\mathbf{v}_k) = \mathbf{p}}{\vdash \{ \mathbf{I}_i \mathbf{v}_i = \text{new } \mathbf{C}_i(\mathbf{v}_{ij}); \mathbf{v}_k! \mathbf{m}(); \} : \mathbf{0} \xrightarrow{\mathbf{t}} \mathbf{p} : \mathbf{m} . \mathbf{G}} \\
 \\
 \text{T-class} \frac{\forall \mathbf{p} \in \mathbf{L}. \exists i. \Delta(\mathbf{f}_i) = \mathbf{p} \quad \Delta, \emptyset \vdash \mathbf{s}_k : \mathbf{L} \upharpoonright^{\mathbf{p}} \mathbf{t} \quad \mathbf{m}_k \text{ is the method of } \mathbf{t} \text{ in } \mathbf{L}}{\Delta \vdash \text{class } \mathbf{C}(\mathbf{I}_i \mathbf{f}_i) \{ \mathbf{T}_j \mathbf{f}_j = \mathbf{e}_j; \mathbf{T}_k \mathbf{m}_k(\mathbf{T}_{kl} \mathbf{v}_{kl}) \{ \mathbf{s}_k \} \} : \mathbf{L}} \\
 \\
 \text{T-}\leq \frac{\Delta, \Gamma \vdash \mathbf{s} : \mathbf{L} \quad \mathbf{L} \equiv \widehat{\mathbf{L}}' \leq \widehat{\mathbf{L}}}{\Delta, \Gamma \vdash \mathbf{s} : \widehat{\mathbf{L}}} \quad \text{T-}; \frac{\Delta, \Gamma \vdash \mathbf{s}_2 : \mathbf{L} \quad \text{inter}(\mathbf{s}_1, \Gamma)}{\Delta, \Gamma \vdash \mathbf{s}_1; \mathbf{s}_2 : \mathbf{L}} \\
 \\
 \text{T-return} \frac{\text{constr}(\mathbf{e}) = \mathbf{C}}{\Delta, \Gamma \vdash \text{return } \mathbf{e}; : \text{Put } \mathbf{t}(\mathbf{C})} \quad \text{T-skip} \frac{}{\Delta, \Gamma \vdash \text{skip} : \text{skip}} \\
 \\
 \text{T-if} \frac{\Delta, \Gamma \vdash \mathbf{s}_1; \mathbf{s}_3 : \mathbf{L} \quad \Delta, \Gamma \vdash \mathbf{s}_2; \mathbf{s}_3 : \mathbf{L}}{\Delta, \Gamma \vdash \text{if}(\mathbf{e}) \{ \mathbf{s}_1 \} \text{else} \{ \mathbf{s}_2 \} \mathbf{s}_3 : \mathbf{L}} \quad \text{T-while} \frac{\Delta, \widetilde{\Gamma} \vdash \mathbf{s}_1 : \mathbf{L}_1 \quad \Delta, \widetilde{\Gamma} \vdash \mathbf{s}_2 : \mathbf{L}_2}{\Delta, \Gamma \vdash \text{while}(\mathbf{e}) \{ \mathbf{s}_1 \} \mathbf{s}_2 : (\mathbf{L}_1)^*. \mathbf{L}_2} \\
 \\
 \text{T-await} \frac{\Delta, \Gamma^{\mathfrak{A}} \vdash \mathbf{s} : \mathbf{L} \quad \Gamma(\mathbf{t}') = (\mathbf{e}, _)}{\Delta, \Gamma \vdash \text{await } \mathbf{e}; \mathbf{s} : \text{Susp}(\mathbf{t}, \mathbf{t}'). \mathbf{L}} \\
 \\
 \text{T-get} \frac{\Gamma(\mathbf{t}) = (\mathbf{e}_2, _) \quad \Delta, \Gamma[\mathbf{t} \mapsto (\mathbf{e}_2, \mathbf{e}_1)] \vdash \mathbf{s} : \mathbf{L} \quad \mathbf{e}_1 \notin \text{im}\Gamma}{\Delta, \Gamma \vdash \mathbf{e}_1 = \mathbf{e}_2. \text{get}; \mathbf{s} : \text{Get } \mathbf{t}. \mathbf{L}} \\
 \\
 \text{T-!} \frac{\Delta, \Gamma[\mathbf{t} \mapsto (\mathbf{e}_1, \perp)] \vdash \mathbf{s} : \mathbf{L} \quad \Delta(\mathbf{e}_2) = \mathbf{p} \quad \mathbf{e}_1 \notin \text{im}\Gamma}{\Delta, \Gamma \vdash \mathbf{e}_1 = \mathbf{e}_2! \mathbf{m}(\mathbf{e}); \mathbf{s} : \mathbf{p}! \mathbf{m}. \mathbf{L}} \\
 \\
 \text{T-case} \frac{\mathbf{C}_i = \mathbf{C}_j \rightarrow \Delta, \Gamma \vdash \mathbf{s}_i : \mathbf{L}_j. \mathbf{L} \quad \forall j. \exists i. \mathbf{C}_i = \mathbf{C}_j \quad \Gamma(\mathbf{t}) = (_, \mathbf{e})}{\Delta, \Gamma \vdash \text{case}(\mathbf{e}) \{ \mathbf{C}_i \Rightarrow \mathbf{s}_i \}_{i \in I} \mathbf{s} : \&_{\mathbf{t}} \{ \mathbf{C}_j : \mathbf{L}_j \}. \mathbf{L}}
 \end{array}$$

■ **Figure 7** Static Type System

480 Rule T- \leq is used for structural congruence and sub-typing. The construction of a syntax-
 481 directed variant of the type system without a special rule for subtyping is standard. Rule T-;
 482 drops a prefix that performs no communication and modifies no location stored in Γ . Rule
 483 T-return checks that the sole remaining action is a Put action and that the correct constructor

is returned. Rule T-**skip** closes the proof if the empty program **skip** is left and no further action is required. This is needed to typecheck loop bodies, where we can always add **skip** at the end. Rule T-**if** splits the derivation into two branches. The type is not changed. Rule T-**while** checks a loop against the Kleene star. The context $\tilde{\Gamma}$ removes all fields and variables modified in the loop body. Rule T-**await** checks that the correct future is synchronized on and removes all fields from the context. Rule T-**get** checks that the correct future is read and stores the information where the read value is available in the context. It ensures that no relevant read value or future is overwritten. Rule T-**!** checks that the correct method on the correct role is called and stores the information where the future is available in the context. It ensures that no other relevant read value or future is overwritten. Rule T-**case** checks a **case** statement against a passive choice by mapping each branch of the statement against some branch of the type. It is ensured that for every specified choice an implemented branch exists and that the read value is indeed stemming from the future containing the choice.

A rule for assignments to copy futures or their read values is easily added, but requires to keep track of a pair of sets of locations and, for simplicity, we refrain from introducing this.

3.4 Globally Dynamic


The globally dynamic part consists of two steps: first, we translate an object-local type into a session automaton, then we translate the session automaton into a user-defined scheduler.

3.4.1 Automaton Extraction


The structure of the translation follows the standard translation of regular expressions into finite automata.

► **Definition 13.** Let \mathbf{L} be a projected object-local type with k tracked futures. Let $\text{pos}(\mathbf{t})$ be the register assigned to \mathbf{t} . The translation of \mathbf{L} into a k -register session automaton is denoted $\mathcal{A}(\mathbf{L})$ and defined as follows:

■ A receiving type $\mathbf{p}^?_{\mathbf{t}\mathbf{m}}$ is translated into an automaton with two states and a single transition that reads $\text{invREv}_{\mathbf{m}}$ and stores the read future in $\text{pos}(\mathbf{t})$:



■ A reactivation type $\text{React}(\mathbf{t})$ is translated into an automaton with two states and a single transition that reads condREv and matches the read future with the one stored in $\text{pos}(\mathbf{t})$.




■ Branching, sequence and repetition are the standard translations of alternative, concatenation and Kleene star into finite automata.

After this construction, standard ϵ -transition elimination is performed.

► **Example 14.** Consider the following scheduling type [29]:

$$\mathbf{L} = (\mathbf{p}^?_{\mathbf{t}_0\mathbf{m}_0} \cdot \mathbf{p}^?_{\mathbf{t}_1\mathbf{m}_1} \cdot \text{React}(\mathbf{t}_0))^*$$

Its translation $\mathcal{A}(\mathbf{L})$ is as follows (the translation yields an ϵ -transition from state 4 to state 1, which is eliminated to give the depicted automaton):



521

522 For formal soundness arguments, we again refer to [20]. Intuitively, the extraction is
 523 sound because the language accepted by the automaton is the same language as the one
 524 generated by the object-local type. Not every extracted session automaton is deterministic,
 525 because the input object-local type may not be deterministic, for example:

$$526 \quad \&_t \left\{ \begin{array}{l} \mathbf{p}^?_m \\ \mathbf{p}^?_m . \mathbf{q}^?_n \end{array} \right\}$$

527 After receiving a call on \mathbf{m} , this type cannot predict which branch to take. We do allow
 528 non-deterministic schedulers, but the implementation issues a warning. A simple syntax
 529 check on the automaton can exclude them.

530 3.4.2 Translation and Integration

531 Given a session automaton, we can finally extract a user-defined scheduler and add instru-
 532 mentation code to ensure correctness.

533 ► **Definition 15.** Let \mathbf{C} be a class that is checked against an object-local type that is transformed
 534 to a scheduling type \mathbf{L} . The instrumented class $\mathbf{C}^{\mathcal{I}}$ is constructed as follows:

- 535 ■ We add a field `Int q = 0`; that models the current state of the scheduling automaton.
- 536 ■ For each register r_i we add a field “`Maybe<Fut<Any>> ri = Nothing`;”.
- 537 ■ The scheduler is as in Def. 16.
- 538 ■ For each method \mathbf{m} we collect all transitions $(q_i, (\text{invREv}, \mathbf{m}), q_{i'})_{i \in I}$ with written register
 539 $\text{reg}(i)$ and add the following as the first statement of \mathbf{m} :

```

540   case this.q {
      qi1 => this.rreg(i1) = Just(destiny); this.q = qi1';
      ⋮
      qim => this.rreg(im) = Just(destiny); this.q = qim';
    }

```

541 This statement saves its future in the given register and updates the automaton state. The
 542 generated scheduler ensures that no default branch is needed.

- 543 ■ For each class \mathbf{C} we collect all transitions $(q_i, (\text{condREv}), q_{i'})_{i \in I}$ with read register $\text{reg}(i)$
 544 and add the following as the first statement after each `await` statement in any method:

```

545   case this.q { qi1 => this.q = qi1'; ... qim => this.q = qim'; }

```

546 Again, the generated scheduler ensures that no default branch is needed and the registers
 547 do not need to be checked against `destiny`.

548 ► **Definition 16.** The generated scheduler ensures that the initializing method with the
 549 hidden name `.init()` is always executed first. The function `filter` is one of the higher-order
 550 functions in *ABS* and takes a function of the form `(params) => code` as its first parameter
 551 and a list as its second.

```

def Maybe<Process> scheduler(List<Process> list,
                             Int q,
                             Maybe<Fut<Any>> r1,
                             Maybe<Fut<Any>> r2) =
  if ( filter((Process p) => method(p) == ".init")(queue) != Nil )
    headOrNothing(filter((Process p) => method(p) == ".init")(queue))
  else case q {
    1 => headOrNothing(
      filter((Process p) => contains(set["publish"],method(p)))(list));
    2 => headOrNothing(
      filter((Process p) => contains(set["request"],method(p)))(list));
  }

```

■ **Figure 8** Scheduler generated from Ex. 11.

```

def Maybe<Process> scheduler(List<Process> list, Int q,
                             Maybe<Fut<Any>> r1, ..., Maybe<Fut<Any>> r_n) =
  if( filter((Process p) => method(p) == ".init")(queue) != Nil )
    headOrNothing(filter((Process p) => method(p) == ".init")(queue))
  else scheduler_body(list, q, r1, ..., r_n);

```

After executing the initializer, the scheduler makes a case distinction over the states $1, \dots, m$ of the scheduling automaton:

```

def Maybe<Process> scheduler_body(List<Process> list, Int q,
                                   Maybe<Fut<Any>> r1, ..., Maybe<Fut<Any>> r_n) =
  case q { 1 => transition1; ... m => transitionm; }

```

The transition transition_i from a state i is modeled as follows: Let m_1, \dots, m_{n_1} be the method names that have outgoing transitions from i labeled with invREv . Let r'_1, \dots, r'_{n_2} be the registers that the read future is compared with in outgoing transitions from i labeled with condREv . The first case checks that the future is allowed and not yet stored, the second case checks that the future is in one of the registers.

```

headOrNothing(filter((Process p) =>
  (contains(set[m1, ..., mn1], method(p)) && !contains(set[r1, ..., rn], destinyOf(p)))
  || contains(set[r'_1, ..., r'_{n_2}], destinyOf(p))
)(list))

```

We return the first process that is in the list and matches, a random scheduler is a straight-forward modification.

► **Example 17.** The (beautified) scheduler generated from Ex. 11 is shown in Fig. 8:

3.5 Soundness and Stateful Session Types

Soundness. Soundness of the type system follows directly from the soundness theorem given for the original, purely static systems [27, 28]:

► **Theorem 18.** Let Prgm be a well-typed ABS program and \mathbf{GP} a global protocol. If $\vdash \text{Prgm} : \mathbf{GP}$ and every object is instrumented with the scheduler type derived by the \vdash relation, then every terminating and non-deadlocking run of Prgm has a trace where the communication events for each object are in the same order as specified in \mathbf{GP} .

Our notion of soundness is not based on subject reduction and progress. Soundness of our system is *only* concerned with protocol adherence, not with deadlock freedom, as discussed above. Adding deadlock checks in session types complicates the system further [28] for little gain, as external tools can be used. We assume that the data types have been checked, so there is no need for a progress theorem. It is similar to *session fidelity* [22], which expresses the same intuition in terms of operational semantics.

Neither do we use a subject reduction theorem. Instead, we give a denotational semantics to session types and regard them as specifications of traces in monadic second-order logic: Any type \mathbf{GP} can be translated into a formula $\mathbb{C}(\mathbf{GP})$ expressing that the communication events for each object are in the same order as specified in \mathbf{GP} . Soundness is then a *model-theoretic* notion that every trace \mathbf{tr} generated by \mathbf{GP} is a model of $\mathbb{C}(\mathbf{GP})$:

$$\vdash \text{Prgm} : \mathbf{GP} \rightarrow \forall \mathbf{tr}. \text{Prgm} \Downarrow \mathbf{tr} \rightarrow \mathbf{tr} \models \mathbb{C}(\mathbf{GP})$$

This model-theoretic treatment of session types allows an elegant connection to symbolic execution and dynamic logic [27] at the cost of an elaborate semantics [14] which we refrain to introduce for space reasons. This semantics is based on merging of local traces, which inhibits us from giving a straightforward subject reduction theorem.

Stateful Session Types So far, our session types do not constrain the execution state or passed data, except the outermost constructor of return values. We implemented an extension of the presented system, where each global call action is annotated with a property. This annotation is preserved during object-local projection and moved to the termination action during projection on a tracked future. Regarding instrumentation, it results in a simple **assert** statement for the dynamic check.

► **Example 19.** We specify that a call of role \mathbf{p} to a method \mathbf{m} results in a postcondition that ensures the return value being larger than field \mathbf{f} :

... $\mathbf{p} \xrightarrow{t} \mathbf{q} : \mathbf{m}(\mathbf{this.f} < \mathbf{result})$...

The return value is saved in a dedicated variable **result** and an assert is added afterwards. If the final statement was “**return** \mathbf{e} ,” before, it now is

Int **result** = \mathbf{e} ; **assert**($\mathbf{this.f} < \mathbf{result}$); **return** **result**;

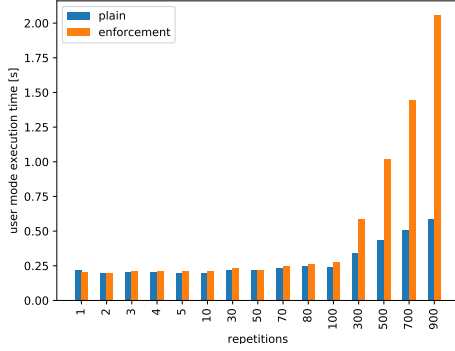
If it depends on the state of the scheduler which postcondition has to be checked, a **case** statement over the possible values of \mathbf{q} is added. This approach is slightly less expressive than other stateful session types for AO [26, 28], but has the benefit that there is no need to translate first-order logic formulas into expressions.

4 Implementation and Evaluation

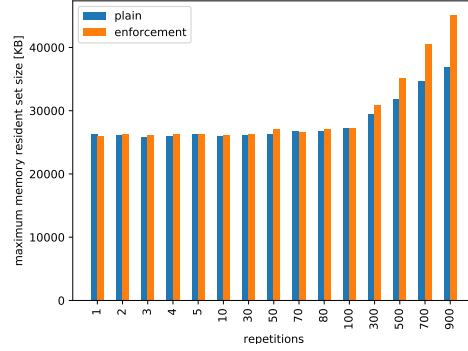
Our system is implemented on top of a slightly modified⁵ version of the ABS compiler [43, 39]. Source code and all examples are accessible at <https://github.com/ahbmr/SessionTypeABS>.

As discussed, we do not handle full ABS and demand that the main block initializes a whole session, each interface plays exactly one role and no objects are created after initialization. The session type is specified in an ASCII variant of Def. 4 in a separate file alongside the other

⁵ Blocking schedulers and access to the future of a process are not yet part of the master branch of ABS.



■ **Figure 9** Execution times of the unmodified (blue) and modified (orange) model for different amounts of repetitions.



■ **Figure 10** Maximum memory resident set size of the unmodified (blue) and modified (orange) model for different ammounts of repetitions.

610 model source files. The **ABS** compiler is used for parsing and (data-)typechecking the input
 611 model. The AST is then used for the static check and enriched with the instrumentation
 612 from the scheduling type. The resulting AST is passed back to the **ABS** compiler, which
 613 parses and typechecks it again.

614 We evaluate the impact of our modifications on the performance of Erlang-based simula-
 615 tions (the standard backend) of **ABS** models. The experiments are performed on a synthetic
 616 benchmark, where one object implementing the role **p** repeatedly calls two methods on
 617 another object of role **q** in a fixed order. The session is specified by the this global type:

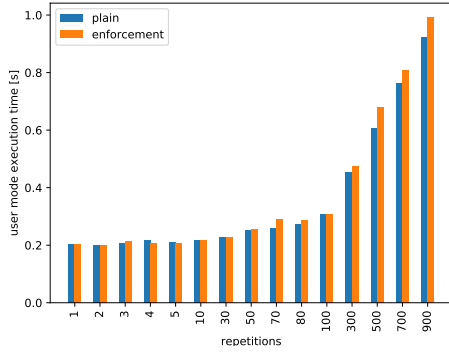
$$618 \quad 0 \xrightarrow{t} \mathbf{p} : \text{init} . \left(\mathbf{p} \xrightarrow{t_{m1}} \mathbf{q} : m1 . \mathbf{q} \downarrow t_{m1} . \mathbf{p} \xrightarrow{t_{m2}} \mathbf{q} : m2 . \mathbf{q} \downarrow t_{m2} \right)^* . \mathbf{p} \downarrow t$$

619 All reported data resulted from executing the model multiple times and averaging the
 620 measurements. Reported *execution times* designate the required run time in user-mode of
 621 the Erlang simulation of a model until termination on a Arch Linux system running Kernel
 622 5.3.7 with a i5-4300U@2.9GHz CPU and 4GiB RAM.

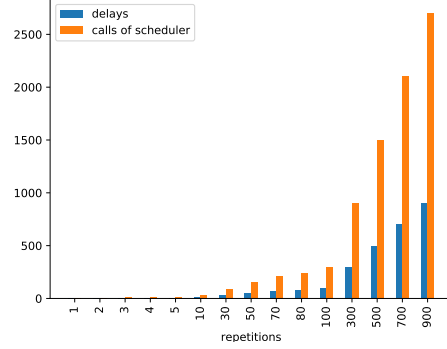
623 *Effect of Increased Object Communication.* By changing the number of times the repeat-
 624 able section of the session type is executed, we observe the behavior of the model simulation
 625 when the number of calls from **p** to **q** increases. We observe that the user-mode execution time
 626 of the simulations is nearly constant and mostly equivalent for the modified and unmodified
 627 version of the model for up to 100 repetitions, see Fig. 9. For higher numbers of repetitions,
 628 execution time increases for both version, but execution time of the modified model grows to
 629 increasing multiples of the execution time of the unmodified one. The maximum memory
 630 resident set size of the Erlang processes develops similarly, see Fig. 10, although the memory
 631 size of the modified model does not grow as rapidly as the execution time.

632 *Comparison to Manual Synchronization.* Instead of letting the generated schedulers
 633 enforce the execution order of methods, we now require **p** and **q** to synchronize every call
 634 by inserting an **await**-statement after each interaction. Even though execution time of the
 635 unmodified and modified model still increases for a high number of repetitions, there is
 636 now little difference between them, see Fig. 11. The overhead of synchronization is roughly
 637 equivalent to or lower than the version relying on the generated schedulers.

638 *Testing the Reordering Capabilities of the Scheduler:* In the previous experiments the
 639 scheduler never delayed activating a process, because there was always one in the queue which



■ **Figure 11** User-mode execution times when using `await` statements. Unmodified model in blue, modified model in orange.



■ **Figure 12** The number of times a scheduler has been invoked (orange) in contrast to the number of times it could not activate any waiting process (blue).

could immediately be scheduled. We now disable static verification, deliberately reverse the calls in the model source and put `duration` statements after each call, causing a delay in the execution.⁶ We do not use synchronization and calls always arrive out of order at `q` and with enough inactivity in between them so that the scheduler of `q` frequently has to delay activation of a process until an acceptable one is available. Here, the modified and unmodified model always complete in almost the same execution time, presumably since the `duration` statements induce enough idle time to contain the overhead of the schedulers. However, we now observe that the scheduler successfully delays and reorders calls, see Fig. 12.

Discussion. A certain overhead must always be expected from instrumentation, but we deem the observed overhead acceptable. The generated schedulers only result in noteworthy overhead when a large number of processes is in the object queue. We conjecture that this effect is mostly an artifact of how the queue is represented for the user-defined scheduler.

5 Related Work

There is a considerable number of papers combining static and dynamic verification, a complete overview is out of scope for this work. We refer to, for example, the introduction of Ahrendt et al. [2] and only review directly related approaches here.

The `StarV00rS` [1, 9] tool combines static and dynamic verification of Java programs as follows: First, it attempts to prove certain properties statically using deductive verification and then it transforms failed proofs into runtime monitors. The static analysis is used to ensure that as little as possible is checked dynamically. `StarV00rS` distinguishes between data and control-flow properties. The static analysis is mainly reducing the need for the computationally heavy data properties (e.g., all values of an array are non-zero) as far as possible, while monitoring control-flow properties can be done statically.

Our approach can be seen from a similar perspective: the object scheduler is handling the control flow inside an object, while the added `assert` statements are handling data properties. The type checker ensures that inside a method, only data properties need to be checked at runtime. It is straightforward to see how the ongoing integration of Session Types into the

⁶ Explicit time behavior is realized in Timed ABS [5] and here only used for evaluation.

667 **Crowbar** prover using Behavioral Program Logic [26] can be used to discard as superfluous
 668 **assert** statements statically.

669 The literature on session types includes approaches that handle protocols as (partially)
 670 dynamic types or mix static and dynamic checks otherwise. The conceptually closest to our
 671 approach is by Bocchi et al. [6], who also use distributed runtime enforcement, but introduce
 672 new components (for example, a queue) to do so. Completely dynamic approaches to session
 673 types are available for the Python language [12] and an actor model [35]. Other, less related,
 674 approaches are:

- 675 ■ Gradual session types [23] transform a dynamically checked *dyadic* session type for
 676 *channels* gradually to a statically checked one during development. The dynamic check
 677 for linearity that is central to gradual session types for channels has no direct counterpart
 678 in our system for AO, because the projection mechanisms differ on a technical level.
- 679 ■ Certain combined approaches, e.g., for Scala [38], draw the line between static and
 680 dynamic by performing only the linearity check at runtime and any other check statically.

681 A further type-based approach is *typestate* [41]. In contrast to session types, it was
 682 developed mainly for OO imperative programs. Typestate models that an object can change
 683 its interface, i.e., the set of exposed methods, over time. This was done statically in the
 684 original work and was subsequently gradualized [42] to combine static and dynamic type
 685 checking. A variant of typestate for concurrent Java, developed by Gerbo & Padovani [17],
 686 dynamically reports violations after injecting monitoring code. The object scheduler in our
 687 approach can be seen as a variant of typestate, but it is *generated*, not specified.

688 Choreographies [8] bear similarity to session types, being global specifications with a
 689 projection mechanism. However, they are mainly used to *generate* code via a correctness-
 690 by-construction approach. This also combines static and dynamic aspects, but reverses the
 691 direction: instead of dynamically ensuring that the static checks are sound, it is statically
 692 ensured (by code generation) that the dynamic behavior is structured correctly. The
 693 distinction between static and dynamic parts becomes even more prominent in the work of
 694 **Gabbrielli** et al. [16, 36, 37], where *dynamic choreographies* are used to generate a dynamic
 695 structure to update the structure of the application or include of new participants.

696 **6 Conclusion**

697 What should be the takeaway message from this work? First, the formalism of session types,
 698 first developed in the context of the π -calculus, and so far mainly used in theoretical investi-
 699 gations, appears in our context as a rather versatile and surprisingly practical specification
 700 mechanism. It is easily conceivable to find a more user-friendly, less mathematical notation
 701 for the global types in Fig. 3 and add IDE support.

702 Second, with the runtime checking approach, session types for AO can form the theoretical
 703 basis for top-down development of *open* distributed systems (with cooperative concurrency).

704 Third, as shown here and in [27], session types integrate well with static checking of
 705 logical properties. The semantic link is a straightforward translation from session types into
 706 logic, while the type systems syntactically ensures to place assertions at suitable locations.

707 *Future Work.* We plan to adopt the **StarV00rS** approach to partially reduce the need
 708 for **assert** statements on method-local level. We are investigating the use of the product line
 709 mechanism of **ABS** [10] to add the monitors, instead of using manual code injection. Using
 710 product lines enables a uniform treatment of code injection in **ABS** and the injection and
 711 removal of runtime monitors at runtime [40]. Furthermore, we plan to investigate the use of
 712 Timed Session Types [34] for Timed **ABS** and Hybrid **ABS** [32].

713 — References —

- 714 1 Wolfgang Ahrendt, Jesús Mauricio Chimento, Gordon J. Pace, and Gerardo Schneider. Verifying data- and control-oriented properties combining static and runtime verification: theory and tools. *Formal Methods Syst. Des.*, 51(1):200–265, 2017.
- 715 2 Wolfgang Ahrendt, Marieke Huisman, Giles Reger, and Kristin Yvonne Rozier. A broader view on verification: From static to runtime and back (track summary). In *ISoLA (2)*, volume 11245 of *Lecture Notes in Computer Science*, pages 3–7. Springer, 2018.
- 716 3 Elvira Albert, Puri Arenas, Antonio Flores-Montoya, Samir Genaim, Miguel Gómez-Zamalloa, Enrique Martin-Martin, German Puebla, and Guillermo Román-Díez. SACO: static analyzer for concurrent objects. In Erika Ábrahám and Klaus Havelund, editors, *TACAS 2014*, volume 8413 of *Lecture Notes in Computer Science*, pages 562–567. Springer, 2014.
- 717 4 Henry G. Baker and Carl E. Hewitt. The incremental garbage collection of processes. In *Proceeding of the Symposium on Artificial Intelligence Programming Languages*, number 12 in SIGPLAN Notices, page 11, August 1977.
- 718 5 Joakim Bjørk, Frank S. de Boer, Einar Broch Johnsen, Rudolf Schlatte, and Silvia Lizeth Tapia Tarifa. User-defined schedulers for real-time concurrent objects. *ISSE*, 9(1):29–43, 2013.
- 719 6 Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *Theor. Comput. Sci.*, 669:33–58, 2017.
- 720 7 Benedikt Bollig, Peter Habermehl, Martin Leucker, and Benjamin Monmege. A fresh approach to learning register automata. In Marie-Pierre Béal and Olivier Carton, editors, *Developments in Language Theory: 17th Intl. Conf. DLT, Marne-la-Vallée, France*, volume 7907 of *LNCS*, pages 118–130. Springer, 2013.
- 721 8 Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centered programming for web services. *ACM Trans. Program. Lang. Syst.*, 34(2):8:1–8:78, 2012.
- 722 9 Jesús Mauricio Chimento, Wolfgang Ahrendt, Gordon J. Pace, and Gerardo Schneider. Starvoors: A tool for combined static and runtime verification of java. In Ezio Bartocci and Rupak Majumdar, editors, *RV 2015*, volume 9333 of *Lecture Notes in Computer Science*, pages 297–305. Springer, 2015.
- 723 10 Dave Clarke, Radu Muschevici, José Proença, Ina Schaefer, and Rudolf Schlatte. Variability modelling in the ABS language. In Bernhard K. Aichernig, Frank S. de Boer, and Marcello M. Bonsangue, editors, *FMCO 2010*, volume 6957 of *Lecture Notes in Computer Science*, pages 204–224. Springer, 2010.
- 724 11 Frank S. de Boer, Vlad Serbanescu, Reiner Hähnle, Ludovic Henrio, Justine Rochas, Crystal Chang Din, Einar Broch Johnsen, Marjan Sirjani, Ehsan Khamespanah, Kiko Fernandez-Reyes, and Albert Mingkun Yang. A survey of active object languages. *ACM Comput. Surv.*, 50(5):76:1–76:39, 2017.
- 725 12 Romain Demangeon, Kohei Honda, Raymond Hu, Romyana Neykova, and Nobuko Yoshida. Practical interruptible conversations: distributed dynamic verification with multiparty session types and python. *Formal Methods Syst. Des.*, 46(3):197–225, 2015.
- 726 13 Crystal Chang Din, Richard Bubel, and Reiner Hähnle. Key-abs: A deductive verification tool for the concurrent modelling language ABS. In Amy P. Felty and Aart Middeldorp, editors, *Automated Deduction - CADE-25 - 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings*, volume 9195 of *Lecture Notes in Computer Science*, pages 517–526. Springer, 2015.
- 727 14 Crystal Chang Din, Reiner Hähnle, Einar Broch Johnsen, Ka I Pun, and Silvia Lizeth Tapia Tarifa. Locally abstract, globally concrete semantics of concurrent programming languages. In *TABLEAUX*, volume 10501 of *Lecture Notes in Computer Science*, pages 22–43. Springer, 2017.
- 728 15 Antonio Flores-Montoya, Elvira Albert, and Samir Genaim. May-happen-in-parallel based deadlock analysis for concurrent objects. In *FMOODS/FORTE*, volume 7892 of *LNCS*, pages 273–288. Springer, 2013.

- 764 16 Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. Guess who's coming:
765 Runtime inclusion of participants in choreographies. In Mário S. Alvim, Kostas Chatzikokolakis,
766 Carlos Olarte, and Frank Valencia, editors, *The Art of Modelling Computational Systems: A
767 Journey from Logic and Concurrency to Security and Privacy - Essays Dedicated to Catuscia
768 Palamidessi on the Occasion of Her 60th Birthday*, volume 11760 of *Lecture Notes in Computer
769 Science*, pages 118–138. Springer, 2019.
- 770 17 Rosita Gerbo and Luca Padovani. Concurrent typestate-oriented programming in java.
771 In Francisco Martins and Dominic Orchard, editors, *Proceedings Programming Language
772 Approaches to Concurrency- and Communication-cEntric Software, PLACES@ETAPS 2019,
773 Prague, Czech Republic, 7th April 2019*, volume 291 of *EPTCS*, pages 24–34, 2019.
- 774 18 Elena Giachino, Cosimo Laneve, and Michael Lienhardt. A framework for deadlock detection
775 in core ABS. *Software and Systems Modeling*, 15(4):1013–1048, 2016.
- 776 19 Dilian Gurov, Reiner Hähnle, and Eduard Kamburjan. Who carries the burden of modularity?
777 In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods,
778 Verification and Validation, 9th Intl. Symp., ISoLA 2020, Rhodes, Greece, LNCS*. Springer,
779 October 2020.
- 780 20 Anton W Haubner. Semi-dynamic session types for ABS. Bachelor thesis, Technical University
781 of Darmstadt, 2019. Available at [https://github.com/ahbnr/SessionTypeABS/blob/master/
782 thesis/thesis_final_pdfa.pdf](https://github.com/ahbnr/SessionTypeABS/blob/master/thesis/thesis_final_pdfa.pdf).
- 783 21 Carl Hewitt, Peter Bishop, and Richard Steiger. A universal modular ACTOR formalism for
784 artificial intelligence. In *Proceedings of the 3rd International Joint Conference on Artificial
785 Intelligence, IJCAI'73*, pages 235–245. Morgan Kaufmann Publishers Inc., 1973.
- 786 22 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types.
787 In George C. Necula and Philip Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-
788 SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco,
789 California, USA, January 7-12, 2008*, pages 273–284. ACM, 2008.
- 790 23 Atsushi Igarashi, Peter Thiemann, Vasco T. Vasconcelos, and Philip Wadler. Gradual session
791 types. *Proc. ACM Program. Lang.*, 1(ICFP):38:1–38:28, 2017.
- 792 24 Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS:
793 A core language for abstract behavioral specification. In Bernhard K. Aichernig, Frank S.
794 de Boer, and Marcello M. Bonsangue, editors, *FMCO 2010*, volume 6957 of *Lecture Notes in
795 Computer Science*, pages 142–164. Springer, 2010.
- 796 25 Eduard Kamburjan. Detecting deadlocks in formal system models with condition synchroniza-
797 tion. *ECEASST*, 76, 2018.
- 798 26 Eduard Kamburjan. Behavioral program logic. In *TABLEAUX*, volume 11714 of *Lecture
799 Notes in Computer Science*, pages 391–408. Springer, 2019.
- 800 27 Eduard Kamburjan. *Modular Verification of a Modular Specification: Behavioral Types as
801 Program Logics*. PhD thesis, Technische Universität Darmstadt, 2020.
- 802 28 Eduard Kamburjan and Tzu-Chun Chen. Stateful behavioral types for active objects. In
803 Carlo A. Furia and Kirsten Winter, editors, *iFM 2018*, volume 11023 of *Lecture Notes in
804 Computer Science*, pages 214–235. Springer, 2018.
- 805 29 Eduard Kamburjan, Crystal Chang Din, and Tzu-Chun Chen. Session-based compositional
806 analysis for actor-based languages using futures. In Kazuhiro Ogata, Mark Lawford, and
807 Shaoying Liu, editors, *ICFEM 2016*, volume 10009 of *Lecture Notes in Computer Science*,
808 pages 296–312, 2016.
- 809 30 Eduard Kamburjan, Crystal Chang Din, Reiner Hähnle, and Einar Broch Johnsen. Asyn-
810 chronous cooperative contracts for cooperative scheduling. In *SEFM*, volume 11724 of *Lecture
811 Notes in Computer Science*, pages 48–66. Springer, 2019.
- 812 31 Eduard Kamburjan, Reiner Hähnle, and Sebastian Schön. Formal modeling and analysis of
813 railway operations with active objects. *Sci. Comput. Program.*, 166:167–193, 2018.
- 814 32 Eduard Kamburjan, Stefan Mitsch, Martina Kettenbach, and Reiner Hähnle. Modeling and
815 verifying cyber-physical systems with hybrid active objects. *CoRR*, abs/1906.05704, 2019.

- 816 **33** Michael Kaminski and Nissim Francez. Finite-memory automata. *Theor. Comput. Sci.*,
817 134(2):329–363, 1994.
- 818 **34** Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. Timed runtime monitoring for
819 multiparty conversations. *Formal Asp. Comput.*, 29(5):877–910, 2017.
- 820 **35** Rumyana Neykova and Nobuko Yoshida. Multiparty session actors. *Logical Methods in*
821 *Computer Science*, 13(1), 2017.
- 822 **36** Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro.
823 Dynamic choreographies - safe runtime updates of distributed applications. In Tom Holvoet
824 and Mirko Viroli, editors, *COORDINATION 2015*, volume 9037 of *Lecture Notes in Computer*
825 *Science*, pages 67–82. Springer, 2015.
- 826 **37** Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro.
827 Dynamic choreographies: Theory and implementation. *Log. Methods Comput. Sci.*, 13(2),
828 2017.
- 829 **38** Alceste Scalas and Nobuko Yoshida. Lightweight session programming in scala. In Shriram
830 Krishnamurthi and Benjamin S. Lerner, editors, *30th European Conference on Object-Oriented*
831 *Programming, ECOOP 2016, July 18-22, 2016, Rome, Italy*, volume 56 of *LIPICs*, pages
832 21:1–21:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- 833 **39** Rudolf Schlatte and abstools Contributors. Modified branch of the abstools compiler version
834 1.8.1 - github source repository. <https://github.com/ahbmr/abstools/tree/thisDestiny>.
835 Accessed: 2019-10-29.
- 836 **40** Rudolf Schlatte, Einar Broch Johnsen, Jacopo Mauro, Silvia Lizeth Tapia Tarifa, and In-
837 grid Chieh Yu. Release the beasts: When formal methods meet real world data. In Frank S.
838 de Boer, Marcello M. Bonsangue, and Jan Rutten, editors, *It's All About Coordination - Essays*
839 *to Celebrate the Lifelong Scientific Achievements of Farhad Arbab*, volume 10865 of *Lecture*
840 *Notes in Computer Science*, pages 107–121. Springer, 2018.
- 841 **41** Robert E. Strom and Shaula Yemini. Typestate: A programming language concept for
842 enhancing software reliability. *IEEE Trans. Software Eng.*, 12(1):157–171, 1986.
- 843 **42** Roger Wolff, Ronald Garcia, Éric Tanter, and Jonathan Aldrich. Gradual typestate. In Mira
844 Mezini, editor, *ECOOP 2011 - Object-Oriented Programming - 25th European Conference,*
845 *Lancaster, UK, July 25-29, 2011 Proceedings*, volume 6813 of *Lecture Notes in Computer*
846 *Science*, pages 459–483. Springer, 2011.
- 847 **43** Peter Y. H. Wong, Elvira Albert, Radu Muschevici, José Proença, Jan Schäfer, and Rudolf
848 Schlatte. The ABS tool suite: modelling, executing and analysing distributed adaptable
849 object-oriented systems. *STTT*, 14(5):567–588, 2012.