

Amalgamating the Session Types and the Object Oriented Programming Paradigms

Sophia Drossopoulou¹, Mariangiola Dezani-Ciancaglini², Mario Coppo²

¹ Department of Computing, Imperial College London

² Dipartimento di Informatica, Università di Torino

Abstract. We suggest an amalgamation of the session type and the object oriented paradigm whereby sessions are amalgamated with methods; where threads consist of the execution of session bodies on objects and communicate with each other through asynchronously sending/receiving objects on channels; where the choice on how to respond to a session request is based on the name of the request and the class of the object receiving the request; where the choice on how to continue a session is made on the basis of the class of the object sent/received; and where sessions are not first class, but can be delegated to other sessions.

We demonstrate our ideas through a small language, *STOOP*, and an example. We formalize a smaller calculus, *Feather^{STOOP}*, and give a formal definition, and prove subject reduction and progress. The latter property is notoriously difficult and sometimes impossible to achieve in sessions languages, however it holds in *Feather^{STOOP}*.

1 Introduction

Session types [14, 18] or channel contracts [9] support lightweight descriptions of communication protocols, by giving types to communication channels, in terms of the types of values sent or received, *e.g.*, the type `?int.!bool` expresses that an integer will be received and then a boolean value will be sent. A session involves channels of dual session type, thus guaranteeing, that after a session has started, the values sent and received will be of the appropriate type.

Session types have been introduced into several different settings, *i.e.*, for variants of the π -calculus [1, 11, 13, 14, 18], for CORBA [19], for functional languages [12, 20], for boxed ambients [10], for the W3C standard description language for Web Services CDL [2, 3, 15, 17, 21], and for object oriented programming languages [4, 6–8, 19]. In the Singularity OS [9] C# has been extended with channel communication across threads which is governed by channel contracts, and verified against state machine descriptions.

Applications of session types span from web services [17] to operating systems [9]. Such applications are often written on object oriented languages, thus motivating research into the *combination* of session types with object oriented languages.

In the current paper we explore the *amalgamation* of the paradigm of session types with that of object oriented programming. This contrasts to earlier work [4, 6–8, 19] which *extended* the oo paradigm, or even existing languages, with features from session

	“traditional” session	“traditional” method	“amalgamated” session/method
request on	a thread	an object	an object
execution starts when	threads reach certain point	immediately	immediately
executed body is	rest of thread	determined by class of receiver	determined by class of receiver
execution	concurrent	same thread	concurrent
communication	any direction interleaved with computations	n-inputs then computation then one output	any direction interleaved with computations

Fig. 1. “traditional” sessions, “traditional” methods, and “amalgamated” sessions/methods .

types. The need for such an amalgamation stems from the observation, that sessions and methods have related, but different features.

We call our approach ST00P (for *Session Types and Object Oriented Programming*). ST00P is only a “language kernel”, since it is only concerned with the amalgamation of the oo and the sessions paradigm, but is agnostic about issues to do with synchronization, distribution, copying of values across local heaps *e.t.c.* ST00P drives the amalgamation very deeply, by unifying sessions and methods, and by basing choices (*i.e.*, which session body to choose, and how to continue with the session) on the class objects (*i.e.*, the receiving object, or the object being sent or received), rather than, on labels, as in traditional session types. ST00P also supports delegation of the current active channel, but does not support higher order sessions, which can be found in most session types systems, *e.g.*, [1, 3, 4, 6, 7, 11, 13, 14, 16, 18]. Furthermore, ST00P supports recursive types, as in traditional session type systems. Because ST00P supports delegation, but not higher order sessions, it satisfies the progress property, a property which is *not* guaranteed in most sessions types systems, even though it is guaranteed in [6–8].

We believe that the amalgamation we suggest leads to a clearer language design, and that ST00P can elegantly express most useful protocols, in a simple way, while retaining the highly desirable progress property.

The current paper is organized as follows: In section 2 we discuss the philosophy of ST00P, in section 3 we describe ST00P in terms of an example, and in section 4 we discuss the approach, how it would be incorporated into a “real” language. We study the language ST00P, in terms of a featherweight, smaller language, $Feather^{ST00P}$. In section 5 we give an overview of $Feather^{ST00P}$, and discuss how it can encode the features used in the example, and how traditional methods can be encoded by ST00P sessions. In sections 6, 7, 8 we describe the syntax, operational semantics and static typing of $Feather^{ST00P}$, and in section 9 we outline the proofs of subject reduction and of progress. In section 10 we draw conclusions and discuss further work.

2 The philosophy of STOOP

In figure 1 we compare “traditional” sessions [7, 14, 20] and “traditional” methods from object oriented languages, with the “amalgamated” session/methods as in STOOP. “Traditional” sessions are invoked on threads in a manner similar to the Ada rendez-vous, and execution starts when two threads reach a certain point in their execution, where they can “service” the session. The body is determined by the rest of the expression currently run by the thread, and is executed concurrently with the requesting thread. “Traditional” sessions allow communication of any number of values in any direction. On the other hand, “traditional” methods are invoked on an object, the body to run is determined by the class of the receiving object, execution is immediate, and takes place in the same thread as the request, and it supports any number of inputs, followed by computation, followed by one output.

In STOOP we have “amalgamated” sessions/methods, which, for brevity, we shall call sessions from now on. Invocation takes place on an object (*e.g.*, a customer asks to withdraw money from a particular ATM machine), and execution takes place immediately, and concurrently with the requesting thread. The body is determined by the class of the receiving object (*e.g.*, the ATM services the withdrawal request according to its class), and any number of communications interleaved with computation are possible.

We believe that the above amalgamated model of session naturally reflects our intuition of services. Furthermore, it can neatly encode “standard” methods.

We now list the features of STOOP :

- Classes have fields and session bodies. Session bodies are selected based on the object classes.
- STOOP is multithreaded, and communication is asynchronous.
- At every step, in each thread, there is one single active channel on which communications are performed.
- A thread may make a session request through $e.s\{e'\}$; then, e is evaluated to an object, and the expression from the session body in its class is executed concurrently with e' , introducing a new pair of fresh channels k and \tilde{k} (one for each communication direction) to perform communications between the session body and e' ³.
- The expressions `send(e)` and `receive` send/receive objects on the active thread channel.
- The expression `sendCase(e) {C1 ▷ e1 || ... || Cn ▷ en}` evaluates object e and sends it on the active thread channel, and then continues with e_i , where C_i is the class that best fits the class of the object sent. The meaning of `receiveCase(x) {C1 ▷ e1 || ... || Cn ▷ en}` is the obvious one⁴.
- `sendWhile(e) {C1 ▷ e1 || ... || Cn ▷ en}` is similar to `sendCase(e) {C1 ▷ e1 || ... || Cn ▷ en}`, except that it allows for enclosed `continue`. Thus, `receiveWhile(x) {C1 ▷ e1 || ... || Cn ▷ en}` has the obvious meaning.

³ k and \tilde{k} plays for channels a role similar to that of `this` for objects.

⁴ Now, choice is based on the class of the object received.

```

1 sessiontype Shopping_ST =
2   !Item.?Money. $\mu\alpha$ .!{ OK  $\triangleright$  !AcctNr.!Money.?( OK  $\triangleright$  ?Date,
3                                     NoMoney  $\triangleright$   $\epsilon$  ) ,
4                                     NoDeal  $\triangleright$   $\epsilon$  ,
5                                     MakeAnOffer  $\triangleright$  ? { Money  $\triangleright$   $\alpha$ ,
6                                                         NoDeal  $\triangleright$   $\epsilon$  } }
7
8 sessiontype Sell_ST =
9   ?Item.!Money. $\mu\alpha$ .?( OK  $\triangleright$  !{ OK  $\triangleright$  ?AcctNr.?Money.!{ OK $\triangleright$ !Date,
10                                                         NoMoney  $\triangleright$   $\epsilon$  } },
11   NoDeal  $\triangleright$   $\epsilon$ ,
12   MakeAnOffer  $\triangleright$  ! { Money  $\triangleright$   $\alpha$ ,
13                       NoDeal  $\triangleright$   $\epsilon$  } }
14
15 sessiontype CreditCheck_ST = ?AcctNr.?Money.!{OK  $\triangleright$  !Date,
16                                     NoMoney  $\triangleright$   $\epsilon$ }

```

Fig. 2. Session types for shopping, for selling, and for checking balance

- The expression `continue` continues execution at the nearest enclosing `sendWhile` or `receiveWhile`.
- The expression `e • s{}` evaluates the object `e`, and delegates to it the current session. The body of the session `s` in the class of that object is executed concurrently, using the current session. At the end, the final value of the body is passed to the current thread.

3 An Example

We adapt a popular example involving a buyer, a seller and a bank. The buyer negotiates a price from a seller, and if and when they have reached agreement, he sends his bank account number so that it gets verified that he has enough money. If he has enough money, he receives the delivery date, otherwise the deal falls through. The seller *delegates* to a bank the part of the session that checks the money in the account. Such delegation has traditionally been expressed through higher order sessions; instead, STOMP can delegate the current session.

The negotiation allows several rounds: the buyer may either accept the price, or break the negotiation, or require a better deal; in the latter case, the seller might respond with the better price, or might break the negotiation. In other words, the negotiation may be broken (and thus the iteration might be exited) at two different places in the protocol, which belong to different parties. Therefore, regular expressions as in [7] are insufficient to express the protocol; instead we need recursive types.

The session type `Shopping_ST`, from figure 2, describes the above protocol from the point of view of the buyer. The part `!Item.?Money` indicates sending an `Item` followed by receipt of a `Money`. The recursive type `$\mu\alpha$.! { OK \triangleright ..., NoDeal \triangleright ...,`

```

1 \begin{lstlisting}
2 class Buyer {
3     Item prodId; AcctNr acctnr; Date delivDate; Seller seller;
4
5     Object Shopping_ST shopping
6     { prodId = ....;
7       seller.sell{
8         send( prodID );
9         Money price = receive;
10        Response resp = ... price ...;
11        sendWhile( resp ){
12          OK ▷ { send(acctnr); send(price);
13              receiveCase( x ) { OK ▷ delivDate=receive; []
14                              NoMoney ▷ null; } }[]
15
16          NoDeal ▷ null; []
17          MakeAnOffer ▷ {
18            receiveCase( x ) { Money ▷ { resp = ... x ...;
19                                    continue; } []
20                                NoDeal ▷ null; } }
21          }
22        }
23    }

```

Fig. 3. The class Buyer

`MakeAnOffer ▷ ... }` describes the negotiation part, whereby an object is sent, and then, depending on whether the actual object sent belongs to class `OK`, `NoDeal`, or `MakeAnOffer`, the first, second or third branch is taken. In the first branch, the account number and the price is sent; then, either `OK` followed by a `Date`, or a `NoMoney` is received. In the third branch, a further object is received, and if that object is a `Money`, then the negotiations resume on the basis of it, whereas if it is a `NoDeal`, the negotiation ends.

In figure 3 we show the implementation of the class `Buyer`. It has the fields `prodId`, `acctnr`, `delivDate`, and `seller`, which (will) contain the identity of the product to be bought, the account number, the delivery date, and the seller.

The class supports one session, called `shopping`, with session type `Shopping_ST` and return type `Object`. In the body of that session the desirable product is determined and stored in `prodId` (line 6). Then, a session request is made to the `seller` to run session `sell` (line 7). Thus, the seller will run the body of `sell` in parallel with the remaining part of the session body of `shopping`, and a connection will be created between the two threads. On this connection, the `Buyer` will send an `Item` (line 8), receive a `Money` and store it in `price` (line 9), and based on its value will calculate his response based on the `price` and store it in `resp` (line 10). On line 11, the buyer enters a loop with `sendWhile`, where he sends `resp`, and branches according to its class. If `resp` is `OK`, indicating acceptance of the price, then the buyer will send his account

```

1 class Seller {
2   Bank bank; Item prodId; BetterPriceOrDeal resp;
3
4   Object Sell_ST sell
5     { String prodID = receive;
6       Money price = ... prodId ...;
7       send(price);
8       receiveWhile( x ){
9         OK ▷ { case( bank●check{ } )
10              { OK ▷ send(... calc.del.date...); []
11                NoMoney ▷ null; } } []
12        NoDeal ▷ null; []
13        MakeAnOffer ▷ { resp = ....price .... ;
14                       sendCase { Money ▷ continue ; []
15                                NoDeal ▷ null; } }
16      }
17    }
18  }

```

Fig. 4. The class Seller

number, and price (line 12); and will receive an object which may be OK, in which case he will receive a Date and store it in `delivDate` (line 13), or will receive a NoMoney (line 14). If `resp` is NoDeal, indicating that the price is unacceptable, then the session terminates. If the response is MakeAnOffer, inviting the seller to make a better offer, then the rest depends on the other party's response. Line 17 contains `receiveCase` indicating that a value will be received, and the remaining steps will be determined by its class. If the value received is a Money, then it will be stored in `price`, and a response will be calculated, and stored in `resp` (line 17) and the recursion will continue (line 18). If the value received is NoDeal, then the loop will be abandoned.

The session type `Sell_ST` from figure 2 describes the above protocol from the point of view of the seller, and is “dual” to `Shopping_ST`. We now consider the class `Seller`, from figure 4. The session body for `sell` starts by receiving the description of an `Item`, calculating and sending its price. Then, in line 8, it enters a `receiveWhile` loop, which is the counterpart to the `sendWhile` loop from `shopping` and performs all the seller's negotiation. The *interesting new feature* shown here is *delegation*, on line 9, whereby, the `bank` is requested to continue the session, using the current connection, and by application of the session body `check`.

The session type for `check` from class `Bank` in figure 5 is the receipt of a `AccntNr` and a `Money` followed by sending either OK, or a `NoMoney` object. Note, that a session body for `check` is *not aware* whether it will be called through a session request, or through delegation. The return type of `check` is the class `OKorNoMoney` which is the common superclass of `OK` and `NoMoney`.

```

1 class Bank {
2     OKorNoMoney CreditCheck_ST check
3     { AccntNr acct = receive;
4       Money amt = receive;
5       response = ... acct ... amt ...
6       sendCase(response){ OK ▷ null; [] NoMoney ▷ null; }
7     }
8 }

```

Fig. 5. The class Bank

4 Discussion of STOOP

As we said in the introduction, STOOP is a “language kernel”, in the sense that it is only concerned with the best amalgamation of the object oriented features with the sessions part, but is agnostic wrt to the remaining features of the language, such as whether the language is distributed or concurrent, and the features for synchronization.

A core STOOP design choice is the amalgamation of session and method. In particular, each session request causes a new thread/process to be spawned. This makes the language semantics very elegant, but, of course, might end up in creating an undesired number of threads/processes. It would be a straightforward extension of the semantics to ensure the execution of “method-like” (*i.e.*, those which first receive values, then compute but do not communicate, and then send a value) sessions in the same thread/process. Slightly more sophisticated schemes would be needed in the distributed setting to ensure that co-located objects execute in the same process.

In STOOP, each object has the capacity to execute at any point in time an unlimited number of threads/processes. Thus an ATM would have the capability to serve any number of customers, even if its cash dispenser were empty. This is clearly an oversimplification, and it would be straightforward to add “readiness”-fields to STOOP so that an object can only execute a certain number of each of its sessions at each point in time, depending also on its internal state. Such schemes can be adapted from those in [5, 9].

So far we have not concerned ourselves with whether we are dealing with a concurrent or a distributed setting. From the point of view of types, and assuming that in the distributed setting all threads would share a class table, the issue is irrelevant. This is why our core language, $Feather^{STOOP}$, is concurrent, thus making the semantics much simpler.

However, when we go to a full language, for a concurrent setting we will need features to ensure synchronization, which we have left to further work. On the other hand, in a distributed setting we might need to consider issues to do with data marshalling, ownership, and heap partitioning, which we have also left to further work.

5 $Feather^{STOOP}$ overview, and encoding of STOOP in $Feather^{STOOP}$

In order to explain and evaluate the STOOP approach, we designed a small, concurrent, imperative language $Feather^{STOOP}$, whose expressions only support the basic oo features

and also session request, session delegation, and branching sending/receiving loops. In more detail, $Feather^{STOOP}$ encompasses the following features:

- classes, inheritance, fields, session bodies,
- field access, sequence, object creation,
- the constructs $receiveW(x)\{\overline{C \triangleright e}\}$ and $sendW(e)\{\overline{C \triangleright e}\}$ which combine receive/send with branching and loops,
- `continue`.

Thus, we omit from $Feather^{STOOP}$ several of the communication constructs used in the example in section 3, *i.e.*, `send` for sending, `sendCase` for sending and branching depending on the class of the object being sent. In $Feather^{STOOP}$ we only include the `sendWhile` construction, because the other two can be encoded in terms of the latter.

In actual fact, the expression $sendWhile(e)\{Obj \triangleright null^{Obj}\}$ encodes $send(e)$. Furthermore, $sendWhile(e)\{\overline{C \triangleright e}\}$ encodes $sendCase(e)\{\overline{C \triangleright e}\}$. In a similar way, `receive`, `receiveCase` can be encoded in terms of `receiveWhile`⁵. Furthermore, we abbreviate `sendWhile` and `receiveWhile` to `sendW` and `receiveW` and `continue` to `cont`.

Finally, we discuss how methods are special cases of $Feather^{STOOP}$ sessions. Thus, the method declaration

$$C \ m \ (C_1 \ x_1, \dots, C_n \ x_n) \{ e \}$$

can be encoded as

$$Obj \ \mu\alpha_1.?\{C_1 \triangleright \dots \mu\alpha_n.?\{C_n \triangleright \mu\alpha.!\{C \triangleright \varepsilon\}\}\dots\} \ m \\ \{receiveW(x_1)\{C_1 \triangleright \{\dots receiveW(x_n)\{C_n \triangleright sendW(e)\{C \triangleright null^{Obj}\}\}\dots\}\}\}$$

Note that the recursion operator μ is irrelevant in this session type, since there is no iteration here.

Similarly, method calls are special cases of session requests. In fact, in a session request of form $o.s\{e\}$, the object o receives the request and s is the session name. Thus, a method call

$$o.m(v_1, \dots, v_n)$$

can be encoded as

$$o.m\{sendW(v_1)\{C_1 \triangleright \dots sendW(v_n)\{C_n \triangleright receiveW(x)\{C \triangleright x\}\}\dots\}\}$$

where C_1, \dots, C_n, C are respectively the parameter and return classes of the method m in the class of the object o .

6 $Feather^{STOOP}$ Syntax

In figure 6 we describe the syntax of $Feather^{STOOP}$, a featherweight representation of our language. We use gray to indicate runtime expressions, *i.e.* expressions that are

⁵ Indeed the encoding of `SendCase` with `SendWhile` in $Feather^{STOOP}$ works only if there are no occurrences of `continue` in the body of a `SendCase` internal to a loop (as it happens, for instance, in the definition of the session `shopping`, see figure 3, line 18). To recover the full `STOOP` expressivity we should add `SendCase` as a separate construct. We avoid doing this since we want to study the basic features of the language, and the treatment of `SendCase` does not introduce any novelty with respect to that of `SendWhile`.

(class)	$D ::= \text{class } C \text{ extends } C \{ \overline{C \text{ f } S} \}$	
(session defn)	$S ::= C \text{ t s } \{ e \}$	
(expression)	$e ::= x \mid v \mid \text{this} \mid e; e \mid$ $\text{new } C \mid e.f := e \mid e.f \mid$ $e.s \{ e \} \mid$ $e \bullet s \{ k \} \mid$ $k.\text{receiveW}(x) \{ \overline{C \triangleright e} \} \mid$ $k.\text{sendW}(e) \{ \overline{C \triangleright e} \} \mid$ $\text{cont} \mid$ $c.\text{receive} \mid c.\text{send}(e)$	standard oo constructs session request session delegation receive, branch and loop send, branch and loop continue run time receive and sent
(value)	$v ::= \text{null}^C \mid o$	
(thread)	$P ::= e \mid P \mid P$	

Fig. 6. $\mathcal{F}eather^{\text{STOOP}}$ Syntax, where syntax occurring only at runtime appears shaded.

produced in the reduction process but are not expected to occur in the source code of a program. We also use the standard convention of denoting by $\bar{\xi}$ a sequence of elements ξ_1, \dots, ξ_n . In the same way, we denote $\#(\bar{\xi})$ the length n of the sequence and with ξ_i the i -th element.

Programs are defined from a collection of classes ranged over by C, D . Each class is represented by a sequence of *fields* of the form $f \ C$, where f represent the field name and C its class and by a list of session expressions of the form $C \text{ t s } \{ e \}$, where C is the return class, t the session type, s the session name, and e the session body. All classes are defined as extensions of the topmost class Obj .

The first line in the definition of expressions describes standard syntax of object oriented languages, except for method call. The successive four lines describe the syntactic constructs which are characteristic of our approach, *i.e.*, the request of a session, the delegation of a session, and the constructs for receiving and sending objects (*communication expressions*).

The body of sendW and receiveW is a sequence of *alternatives*, $\{ \overline{C \triangleright e} \}$, whose choice depends on the class of the value sent or received. The alternatives e_i ending with a cont expression express repetition of the loop, the other ones represent exit points.

Channels are implicit in the source language syntax. At runtime, communication channels k are introduced at each new session request. We indicate with the operation $\tilde{\cdot}$ the dual, where \tilde{k} is again a runtime channel, and where the operation is such that $\tilde{\tilde{k}} = k$. In each thread we use two channels: one for receiving, and its dual for sending. Note that the meaning of polarities is different from that in [11], where polarities simply represent the two ends of a (unique) session channel.

The run time receive and sent expressions $c.\text{receive}$ and $c.\text{send}(e)$ are handy to terminate the session delegation, see rule (SessDel-R). A fresh channel c is created to allow this (unique) communication.

Objects are ranged over by o and parallel threads by P .

7 Feather^{STOOP} Operational Semantics

Field lookup

$$fields(\text{Obj}) = \bullet \quad \frac{fields(C') = \overline{f'D'} \quad \text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT}}{fields(C) = \overline{f'D'}, \overline{fD}}$$

Field type lookup

$$\frac{fD \in fields(C)}{fType(C, f) = D}$$

Session lookup

$$sessions(\text{Obj}) = \bullet \quad \frac{sessions(C') = \overline{S'} \quad \text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT}}{sessions(C) = \overline{S'}, \overline{S}}$$

Session type lookup

$$\frac{\text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT} \quad C'' \text{ t s } \{ \mathbf{e} \} \in \overline{S}}{sType(\mathbf{s}, C) = \mathbf{t}}$$

$$\frac{\text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT} \quad C'' \text{ t s } \{ \mathbf{e} \} \in \overline{S}}{rType(\mathbf{s}, C) = C''}$$

$$\frac{\text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT} \quad \mathbf{s} \notin \overline{S}}{sType(\mathbf{s}, C) = sType(\mathbf{s}, C')} \quad \frac{\text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT} \quad \mathbf{s} \notin \overline{S}}{rType(\mathbf{s}, C) = rType(\mathbf{s}, C')}$$

Session body lookup

$$\frac{\text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT} \quad C'' \text{ t s } \{ \mathbf{e} \} \in \overline{S}}{sBody(\mathbf{s}, C) = \mathbf{e}}$$

$$\frac{\text{class } C \text{ extends } C' \{ \overline{fD\overline{S}} \} \in \text{CT} \quad \mathbf{s} \notin \overline{S}}{sBody(\mathbf{s}, C) = sBody(\mathbf{s}, C')}$$

Fig. 7. Lookup Functions.

We assume a fixed, global class table CT, which as usual contains Obj as topmost class. In figure 7 we define the auxiliary functions for table lookup used in the operational semantics and typing rules.

Objects and values passed in asynchronous communications are stored in a *heap*. A heap h is a finite mapping with domain objects and channel names. Its syntax is given by:

$$h ::= [] \mid o \mapsto (C, \overline{f : \overline{v}}) \mid k \mapsto \overline{v} \mid c \mapsto v \mid c \mapsto () \mid h+h$$

where $+$ denotes heap concatenation.

Standard OO and Thread Reduction

$$\begin{array}{c}
\text{(FLDASS-R)} \\
\hline
o.f := v, h \longrightarrow v, h[o \mapsto h(o)[f \mapsto v]] \\
\hline
\end{array}
\qquad
\begin{array}{c}
\text{(SEQ-R)} \\
\hline
v; e, h \longrightarrow e, h \\
\hline
\end{array}
\qquad
\begin{array}{c}
\text{(FLD-R)} \\
\hline
h(o) = (C, \bar{f} : \bar{v}) \\
\hline
o.f_i, h \longrightarrow v_i, h \\
\hline
\end{array}$$

$$\begin{array}{c}
\text{(NEWC-R)} \\
\hline
\text{fields}(C) = \overline{C\bar{f}} \quad o \notin h \\
\hline
\text{new } C, h \longrightarrow o, h[o \mapsto (C, \bar{f} : \text{null}^C)] \\
\hline
\end{array}
\qquad
\begin{array}{c}
\text{(CONG-R)} \\
\hline
e, h \longrightarrow e', h' \\
\hline
E[e], h \longrightarrow E[e'], h' \\
\hline
\end{array}
\qquad
\begin{array}{c}
\text{(PAR-R)} \\
\hline
e, h \longrightarrow e', h' \\
\hline
P|e, h \longrightarrow P|e', h' \\
\hline
\end{array}$$

Communication Reduction

$$\begin{array}{c}
\text{(SESSREQ-R)} \\
\hline
h(o) = (C, -) \quad sBody(s, C) = e' \quad k, \bar{k} \notin h \\
\hline
E[o.s \{e\}], h \longrightarrow E[e[k]] \mid e'[o/this][\bar{k}], h[k, \bar{k} \mapsto ()] \\
\hline
\end{array}$$

$$\begin{array}{c}
\text{(SESSDEL-R)} \\
\hline
h(o) = (C, -) \quad sBody(s, C) = e' \quad c \notin h \\
\hline
E[o \bullet s \{k\}], h \longrightarrow E[c.receive] \mid c.send(e'[o/this][k]), h[c \mapsto ()] \\
\hline
\end{array}$$

$$\begin{array}{c}
\text{(SENDWHILE-R)} \\
\hline
e, h \longrightarrow^* v, h' \quad h'(\bar{k}) = \bar{v} \quad h'(v) = (C, -) \quad C \Downarrow \bar{C} = C_i \\
\hline
k.sendW(e)\{\bar{C} \triangleright \bar{e}\}, h \longrightarrow e_i[k.sendW(e)\{\bar{C} \triangleright \bar{e}\}/cont], h'[\bar{k} \mapsto \bar{v} :: v] \\
\hline
\end{array}$$

$$\begin{array}{c}
\text{(RECEIVEWHILE-R)} \\
\hline
h(k) = v :: \bar{v} \quad h(v) = (C, -) \quad C \Downarrow \bar{C} = C_i \\
\hline
k.receiveW(x)\{\bar{C} \triangleright \bar{e}\}, h \longrightarrow e_i[v/x][k.receiveW(x)\{\bar{C} \triangleright \bar{e}\}/cont], h[k \mapsto \bar{v}] \\
\hline
\end{array}$$

$$\begin{array}{c}
\text{(SEND-R)} \\
\hline
c.send(v), h \longrightarrow \text{null}^{Obj}, h[c \mapsto v] \\
\hline
\end{array}
\qquad
\begin{array}{c}
\text{(RECEIVE-R)} \\
\hline
h(c) = v \\
\hline
c.receive, h \longrightarrow v, h \setminus c \\
\hline
\end{array}$$

Fig. 8. Expression Reduction.

Objects names are mapped to pairs consisting of the class of the object and the list of its fields associated with their values. Channel names are mapped to queues of values. The notation $h[o \mapsto (C, \bar{f} : \bar{v})]$ (where C is the class name and $\bar{f} : \bar{v}$ denotes the association of field names with their values (in the instance considered)) represents both the update of the object o in h , if o is already in the domain of h , and the addition of the object o to h otherwise. The heap produced by $h[k \mapsto \bar{v}]$ maps the live channel k to the queue \bar{v} with the same convention as before. The channels c created for handling session delegation are instead used to store at most a single value and have no dual.

$$e[k] = \begin{cases} e_1[k] \bullet s \{k\} & \text{if } e = e_1 \bullet s \{ \}, \\ k.\text{receiveW}(x)\{C \triangleright e[k]\} & \text{if } e = \text{receiveW}(x)\{\overline{C \triangleright e}\}, \\ k.\text{sendW}(x)\{C \triangleright e[k]\} & \text{if } e = \text{sendW}(x)\{\overline{C \triangleright e}\}, \\ e_1[k]; e_2[k] & \text{if } e = e_1; e_2, \\ e_1[k].f := e_2[k] & \text{if } e = e_1.f := e_2, \\ e_1[k].f & \text{if } e = e_1.f, \\ e_1[k].s \{e_2\} & \text{if } e = e_1.s \{e_2\}, \\ e & \text{otherwise.} \end{cases}$$

Fig. 9. Channel Replacement.

$$e[e'/\text{cont}] = \begin{cases} e' & \text{if } e = \text{cont}, \\ e_1; e_2[e'/\text{cont}] & \text{if } e = e_1; e_2, \\ e & \text{otherwise.} \end{cases}$$

Fig. 10. Continuation Replacement.

Heap membership for object identifiers and channels is checked using standard set notation, by identifying h with its domain, we can also write $o \in h$, $k \in h$, and $c \in h$. We convene that $h(\text{null}^C) = (C, -)$.

In order to discuss the operational semantics of $\mathit{Feather}^{\text{STOOP}}$ we start by listing the evaluation contexts (based obviously on runtime syntax).

$$E ::= [-] \mid E.f \mid E;e \mid E.f := e \mid o.f := E \mid E.s\{e\} \mid E \bullet s \{k\} \mid c.\text{send}(E)$$

Figure 8 gives the reduction rules of $\mathit{Feather}^{\text{STOOP}}$. Since value are passed and stored in the heap the reduction relation is of the form $e, h \longrightarrow e', h'$. The first six rules define execution of standard object oriented constructions, execution within a context, and execution for threads. In the sixth rule, *i.e.*, rule (Par-R) parallel composition is considered modulo structural equivalence. As usual, we define structural equivalence rules asserting that parallel composition is associative and commutative:

$$P|P_1 \equiv P_1|P \quad P|(P_1|P_2) \equiv (P|P_1)|P_2 \quad P \equiv P' \Rightarrow P|P_1 \equiv P'|P_1$$

The communication rules of figure 8 are more interesting, especially the first four. We make use of the special channel update operation $[...]$, and of the continuation replacement operation $[.../\text{cont}]$, defined in figures 9 and 10. Thus, $e[k]$ is the expression e in which all occurrences of receive, send, and delegation expression which are *not* within the body of a session request are extended so that they explicitly mention

the channel k they will use. Also, $e[e'/\text{cont}]$ is the expression e in which all occurrences of cont that *are not* within the body of a send/receive loop are replaced by e' . Thus, occurrences of cont preserve the correct nested structure of while expressions. The typing rules assure that cont as proper subexpression can only occur as the last expression in some sequential composition (operator “;”) or as an alternative of a while expression (see Proposition 8.1(1)).

The send communication rule puts the value v , *i.e.*, the result of evaluating the expression e in the current thread, in the queue associated to the dual channel \tilde{k} of the communication channel k . The computation then proceeds with the expression $e_i[k.\text{sendW}(e)\{\overline{C \triangleright e}\}/\text{cont}]$, *i.e.*, with the expression e_i where all outermost occurrences of cont are replaced by $k.\text{sendW}(e)\{\overline{C \triangleright e}\}$, if C_i is the smallest class in \overline{C} to which the value v belongs. This is given by the conditions $h'(v) = (C, _)$ and $C \Downarrow \overline{C} = C_i$, where we define:

$$C \Downarrow \overline{C} = \begin{cases} C_i & \text{if } C <: C_i \text{ and } \forall j. C <: C_j \implies C_i <: C_j, \\ \perp & \text{otherwise.} \end{cases}$$

Dually the receive communication rule takes a value v from the queue associated to channel k and returns the expression $e_i[v/x][k.\text{receiveW}(x)\{\overline{C \triangleright e}\}/\text{cont}]$, if C_i is the smallest class in \overline{C} to which the value v belongs.

Notice that in well typed expressions $C \Downarrow \overline{C}$ is always defined: this can be easily proved by inspecting the typing rules.

Note also that $k.\text{sendW}(E)\{\overline{C \triangleright e}\}$ is not a context. This allows us to write a simpler semantics for sendW , since no communication is allowed in E (this is assured by the typing rules). Note also the mixture of small and large step semantics in rule (SendWhile-R).

Rule (SessReq-R) creates a pair of fresh dual channels k, \tilde{k} and replaces the occurrence of the session request $o.s\{e\}$ in the context $E[\]$ by the expression $e[k]$, *i.e.*, the expression e where all the communication and delegation commands of the current session use the channel k . Moreover it generates the new thread $e'[o/\text{this}][\tilde{k}]$, where e' is the body of the session s in the class of the object o . Thus, an object can service *any number* of session requests; this is the object-oriented counterpart to the (conn) rule from [16].

Rule (SessDel-R) creates a fresh channel c which allows to synchronize the end of the execution of the body of the session s opened by a delegation. The session delegation $o \bullet s\{k\}$ is replaced by the expression $c.\text{receive}$, which expects a value via channel c and returns that value (see rule (Receive-R)). Moreover the new thread $c.\text{send}(e'[o/\text{this}][k])$ is generated, where e' is the body of the session s in the class of the object o . This thread first evaluates the expression $e'[o/\text{this}][k]$, then sends the resulting value via channel c to the originating thread and lastly returns null^{obj} (see rule (Send-R)). This allows the delegation of a part of the communication via the channel k to the object o : this delegation is transparent for the thread using the dual channel \tilde{k} . When the delegated job is over the original thread can resume the communication via the channel k . Note that $c.\text{send}(E)$ is an evaluation context, so a delegated session can perform communications and delegations.

8 Feather^{STOOP} Types and Typing Rules

8.1 Types

Session types, τ , describe the communications that take place during a session. The syntax of session types is:

$$\tau ::= \varepsilon \mid \tau_1.\tau_2 \mid \mu\alpha.\dagger\{\overline{C \triangleright \tau}\} \mid \alpha \mid \square$$

where we use \dagger as a convenient abbreviation that ranges over $\{!, ?\}$. By ε we denote the *empty* communication, and the *concatenation* $\tau_1.\tau_2$ expresses the communications in τ_1 followed by those in τ_2 . The session type ε is the neutral element of concatenation, so that $\varepsilon.\tau = \tau = \tau.\varepsilon$ for all τ .

The types $\mu\alpha.!\{\overline{C \triangleright \tau}\}$ and $\mu\alpha.?\{\overline{C \triangleright \tau}\}$ express respectively the sending and the receiving of a value: depending on the class of this value the communication will proceed with one of the $\overline{\tau}$. More precisely if C is the class of the value and $C \Downarrow \overline{C} = C_i$, then the communication will proceed with τ_i . Notice that the variable α can occur in $\overline{\tau}$ with the usual meaning of representing the whole session type.

The type \square is used only to type the command `cont`: it plays the role of a place holder which will be replaced by a type variable when the while expression is completed (see rules (RecW-T) and (SendW-T)).

Notice that the current session types are considerably simpler than those in [7].

8.2 Typing of Channel Free Expressions

Environments and Well-formed Environments

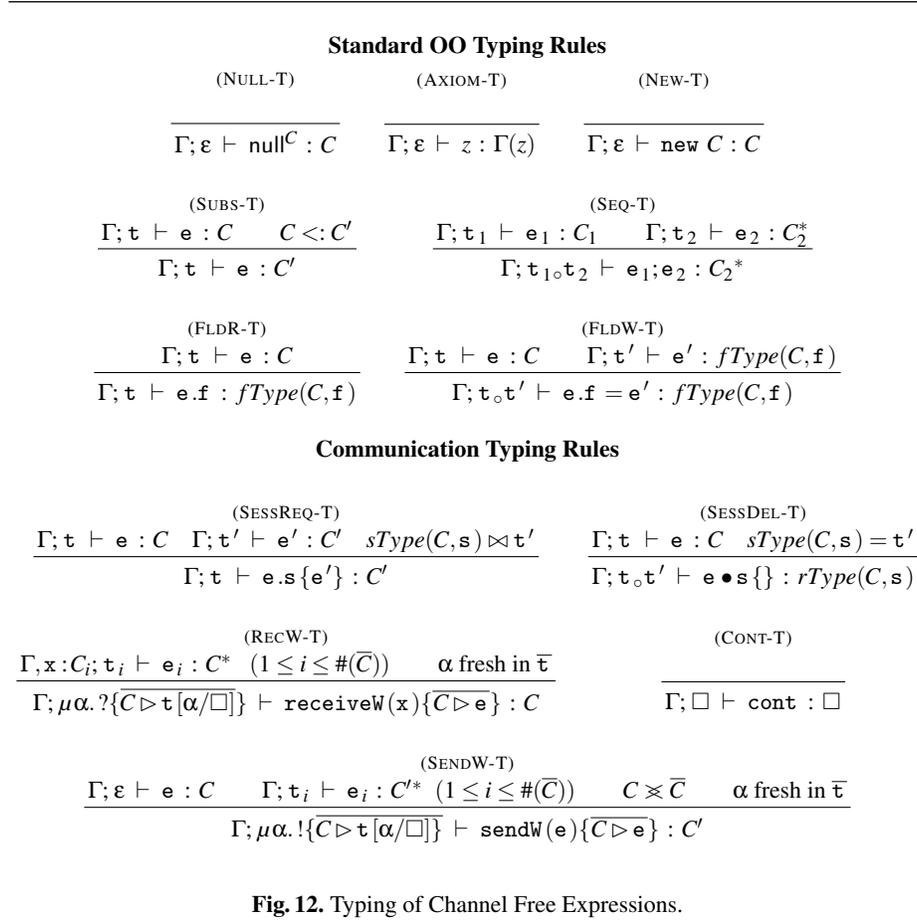
Emp	EVar	EOid	Ethis
$\emptyset \vdash \text{ok}$	$\frac{C \in \text{dom}(\text{CT}) \quad x \notin \text{dom}(\Gamma)}{\Gamma, x : C \vdash \text{ok}}$	$\frac{C \in \text{dom}(\text{CT}) \quad o \notin \text{dom}(\Gamma)}{\Gamma, o : C \vdash \text{ok}}$	$\frac{C \in \text{dom}(\text{CT}) \quad \text{this} \notin \text{dom}(\Gamma)}{\Gamma, \text{this} : C \vdash \text{ok}}$

Subtyping

$$\frac{C \in \text{dom}(\text{CT})}{C <: C} \quad \frac{C <: C' \quad C' <: C''}{C <: C''} \quad \frac{\text{class } C \text{ extends } C' \quad \{\overline{fDS}\} \in \text{CT}}{C <: C'}$$

Fig. 11. Environments and Subtyping.

In this subsection we define typing for user expressions, in which communication channels are implicit. For technical reasons it is useful to consider also expressions with occurrences of object identifiers, which are not directly expressible in user syntax. We



call *channel free expressions* these expressions. The environments therefore will contain also type assignment to object identifiers, this allows a simpler integration of the user expressions typing rules in the runtime typing rules. The typing judgment has the shape

$$\Gamma; \mathbf{t} \vdash \mathbf{e} : C,$$

where the environment Γ maps this, variables and objects to classes, and \mathbf{t} represents the session type of the active channel. Type environments are defined by the syntax:

$$\Gamma ::= \emptyset \mid \Gamma, \mathbf{x} : C \mid \Gamma, \text{this} : C \mid \Gamma, \mathbf{o} : C$$

Figure 11 defines well formed environments and subclassing in the standard way. Figure 12 gives the typing rules. We take the metavariable z to range over this, or a variable x or an object identifier \mathbf{o} . We use of the operation \circ , which is the concatenation,

provided that the first session type does not terminate with \square .

$$\tau \circ \tau' = \begin{cases} \perp, & \text{if } \tau = \tau''.\square, \\ \tau.\tau', & \text{otherwise.} \end{cases}$$

This operation is used for example in rule (Seq-T) to represent that first the communications in e_1 and then those in e_2 are performed.

If C is a class we use C^* to denote either C itself or \square .

To assure a safe communication between two threads we must require that their session types are *dual*, *i.e.*, that each send will correspond to a receive and vice versa. The exchanged values must also be of one of the classes expected by the receiver. Lastly all possible choices on the basis of the class of the exchanged value must continue with session types which are dual of each other. The duality is then the symmetric relation generated by:

$$\begin{array}{c} \varepsilon \bowtie \varepsilon \quad \alpha \bowtie \alpha \quad \frac{\tau_1 \bowtie \tau_1' \quad \tau_2 \bowtie \tau_2'}{\tau_1.\tau_2 \bowtie \tau_1'.\tau_2'} \\ \hline \frac{\bar{C} <: \bar{C}' \quad (C_i \Downarrow \bar{C}' = C'_j \text{ or } C'_j \Downarrow \bar{C} = C_i) \implies \tau_i \bowtie \tau'_j}{\mu\alpha.!\{C \triangleright \tau\} \bowtie \mu\alpha.?\{C' \triangleright \tau'\}} \end{array}$$

where $\bar{C} <: \bar{C}'$ iff $\forall C_i \exists C_j. C_i <: C'_j$, *i.e.*, a sequence of classes is smaller than another one if each class in the first sequence is subclass of a class in the second sequence.

The duality relation is used in rule (SessReq-T) to assure that the body of the session s in class C and the expression e' will properly communicate.

In typing session delegation (rule (SessDel-T)) we take into account that the whole expression will be replaced by the result of the evaluation of the session body (cf. the reduction rule (SessDel-R)).

The rules (RecW-T) and (SendW-T) require that all possible alternative expressions which do not end in `cont` have the same class, but they can implement different communication sequences. Note that the type of `cont` allows to use it only as last statement of a “while” branch (see Proposition 8.1(2)). Note also that it is allowed that all alternatives of a “while” expression end with `cont`, determining a never ending loop. In this case the result type is an arbitrary class. We could easily avoid this by requiring that at least one of the alternatives be typed by a class.

Rule (SendW-T) prescribes that the evaluation of the value to be sent does not involve communications and the *compatibility* of the class C of this value with the classes \bar{C} (notation $C \bowtie \bar{C}$), *i.e.*, that C is a superclass of each class in \bar{C} , and that any subclass of C occurring in the class table is also a subclass of one of the \bar{C} .

$$C \bowtie \bar{C} \text{ iff } \begin{cases} \forall C' \in \bar{C}. C' <: C \\ C' <: C \implies \exists C'' \in \bar{C}. C' <: C''. \end{cases}$$

Figure 13 defines well-formed class tables. Rule **S-ok** type-checks the session bodies with respect to a class C taking as environments the association between this and C .

We conclude this section by showing that in a well formed expression all occurrences of `cont` properly represent loop repetitions.

$$\begin{array}{c}
\text{(S-OK)} \\
\frac{\{\text{this} : C\}; \mathfrak{t} \vdash e : C'}{C' \mathfrak{t} \mathfrak{s} \{e\} : \text{ok in } C} \\
\\
\text{(C-OK)} \\
\frac{\overline{S} : \text{ok in } C}{\text{class } C \text{ extends } C' \{ \overline{f} \overline{D} \overline{S} \} : \text{ok}} \\
\\
\text{(CT-OK)} \\
\frac{\text{class } C \text{ extends } C' \{ \overline{f} \overline{D} \overline{S} \} : \text{ok} \quad \text{CT} : \text{ok}}{\text{CT}, \text{class } C \text{ extends } C' \{ \overline{f} \overline{D} \overline{S} \} : \text{ok}}
\end{array}$$

Fig. 13. Well-formed Class Tables.

- Proposition 8.1.** 1. All occurrences of `cont` as proper subexpressions of well typed expressions are last subexpressions of some sequential compositions or form an alternative of a while expression.
2. All occurrences of `cont` in well typed session requests or session bodies are last subexpressions of some alternatives of while expressions.

Proof. (1) The type of `cont` is \square and only the typing rules for bodies of while expressions and for sequential composition allow the type \square .

(2) The type \square has no dual type and so the conclusion follows from (1) and from the typing rules for while expressions, session body and session request.

8.3 Typing of Runtime Expressions

In typing run time expressions we need to take into account the types of the live channels. For this reason the typing judgements for runtime expressions have the shape

$$\Gamma; \Sigma \Vdash e : C,$$

where Σ denote *session environments* which map live channels to session types and are defined by the following syntax:

$$\Sigma ::= \emptyset \mid \Sigma, k : \mathfrak{t}.$$

We will call *standard environments* the environments Γ .

Figure 14 defines well formed session environments and gives the typing rules for run time expressions, which differ from those for channel free expressions for having session environments instead of an unique session type. For this reason we need to extend the composition of session types to session environments. The *composition of session environments* is the natural generalisation of the composition of session types:

$$\Sigma \circ \Sigma' = \begin{cases} \perp, & \text{if } \exists k \text{ with } \Sigma(k) \circ \Sigma'(k) = \perp \\ \Sigma'', & \text{otherwise, where} \\ & \text{dom}(\Sigma'') = \text{dom}(\Sigma) \cup \text{dom}(\Sigma'), \text{ and} \\ & \Sigma''(k) = \begin{cases} \Sigma(k) \circ \Sigma'(k), & \text{if } k \in \text{dom}(\Sigma) \cap \text{dom}(\Sigma'); \\ \Sigma(k), & \text{if } k \in \text{dom}(\Sigma) \setminus \text{dom}(\Sigma'); \\ \Sigma'(k), & \text{otherwise.} \end{cases} \end{cases}$$

Well formed Session Environments

$$\frac{}{\emptyset \vdash \text{ok}} \quad \text{(SEMP)} \qquad \frac{k \notin \text{dom}(\Sigma)}{\Sigma, k : \text{t ok}} \quad \text{(SERC)}$$

Runtime Typing Rules

$$\begin{array}{c} \text{(NULL-RT)} \qquad \text{(REC-RT)} \qquad \text{(SEND-RT)} \\ \hline \Gamma; \emptyset \Vdash \text{null}^C : C \qquad \Gamma; \emptyset \Vdash \text{c.receive} : C \qquad \Gamma; \Sigma \Vdash \text{c.send}(e) : C \\ \hline \text{(AXIOM-RT)} \qquad \text{(SUBS-RT)} \qquad \text{(SEQ-RT)} \\ \hline \Gamma; \emptyset \Vdash z : \Gamma(z) \qquad \Gamma; \Sigma \Vdash e : C \quad C <: C' \qquad \Gamma; \Sigma_1 \Vdash e_1 : C_1 \quad \Gamma; \Sigma_2 \Vdash e_2 : C_2^* \\ \hline \Gamma; \emptyset \Vdash z : \Gamma(z) \qquad \Gamma; \Sigma \Vdash e : C' \qquad \Gamma; \Sigma_1 \circ \Sigma_2 \Vdash e_1; e_2 : C_2^* \\ \hline \text{(NEW-RT)} \qquad \text{(FLDR-RT)} \qquad \text{(FLDW-RT)} \\ \hline \Gamma; \emptyset \Vdash \text{new } C : C \qquad \Gamma; \Sigma \Vdash e.f : fType(C, f) \qquad \Gamma; \Sigma \Vdash e : C \quad \Gamma; \Sigma' \Vdash e' : fType(C, f) \\ \hline \Gamma; \Sigma \circ \Sigma' \Vdash e.f = e' : fType(C, f) \\ \hline \text{(SESSREQ-RT)} \\ \hline \Gamma; \Sigma \Vdash e : C \quad \Gamma; \mathfrak{t}' \vdash e' : C' \quad sType(C, s) \bowtie \mathfrak{t}' \\ \hline \Gamma; \Sigma \Vdash e.s\{e'\} : C' \\ \hline \text{(SESSDEL-RT)} \\ \hline \Gamma; \Sigma \Vdash e : C \quad sType(C, s) = \mathfrak{t}' \\ \hline \Gamma; \Sigma \circ \{k : \mathfrak{t}'\} \Vdash e \bullet s\{k\} : rType(C, s) \\ \hline \text{(RECW-RT)} \qquad \text{(CONT-T)} \\ \hline \Gamma, x : C_i; \Sigma, k : \mathfrak{t}_i \Vdash e_i : C^* \quad 1 \leq i \leq \#(\bar{C}) \quad \alpha \text{ fresh in } \bar{\mathfrak{t}} \\ \hline \Gamma; \Sigma, k : \mu\alpha. !\{C \triangleright \mathfrak{t}[\alpha/\square]\} \Vdash k.\text{receiveW}(x)\{C \triangleright e\} : C \quad \Gamma; \square \Vdash \text{cont} : \square \\ \hline \text{(SENDW-RT)} \\ \hline \Gamma; \emptyset \Vdash e : C' \quad \Gamma; \Sigma, k : \mathfrak{t}_i \Vdash e_i : C^* \quad 1 \leq i \leq \#(\bar{C}) \quad C' \bowtie \bar{C} \quad \alpha \text{ fresh in } \bar{\mathfrak{t}} \\ \hline \Gamma; \Sigma, k : \mu\alpha. !\{C \triangleright \mathfrak{t}[\alpha/\square]\} \Vdash k.\text{sendW}(e)\{C \triangleright e\} : C \end{array}$$

Fig. 14. Typing of Runtime Expressions.

Notice that in rule (SessReq-T) we are making use of the judgment $\Gamma; \mathfrak{t}' \vdash e' : C'$, where the expression e' does not contain channels, but it can contain object identifiers.

9 Fundamental Properties

In this section we state the fundamental properties that assure that our system is well founded: subject reduction and progress. Subject reduction assures that well typing is preserved during reduction, excluding the possibility of unexpected run time errors (in our case values that could not be handled by the operational semantics). Progress assures that a well typed expression will run until a meaningful value will be produced.

9.1 Subject reduction

A first observation is that a user expression which can be typed in the system of Subsection 8.2 can also be typed in the system of Subsection 8.3 once the current channel is made explicit. The proof by induction on the structure of expressions is standard.

Proposition 9.1. $\Gamma; \tau \vdash e : C$ implies $\Gamma; \{k : \tau\} \Vdash e[k] : C$.

In order to state subject reduction we start by defining a judgment, $\tau \sqsubseteq \tau'$, which describes that a session τ' is at a later stage than another session τ .

$$\begin{array}{c} \text{(LATER-1)} \\ \hline \tau \sqsubseteq \tau \end{array} \quad \begin{array}{c} \text{(LATER-2)} \\ \tau \sqsubseteq \tau'' \quad \tau'' \sqsubseteq \tau' \\ \hline \tau \sqsubseteq \tau' \end{array} \quad \begin{array}{c} \text{(LATER-3)} \\ \mu\alpha. \dagger \{C \triangleright \tau\} \sqsubseteq \tau_i [\mu\alpha. \dagger \{C \triangleright \tau\} / \alpha] \\ \hline \mu\alpha. \dagger \{C \triangleright \tau\} \sqsubseteq \tau_i [\mu\alpha. \dagger \{C \triangleright \tau\} / \alpha] \end{array} \quad \begin{array}{c} \text{(LATER-4)} \\ \tau \sqsubseteq \tau' \\ \hline \tau.\tau'' \sqsubseteq \tau'.\tau'' \end{array}$$

Notice that not all types which are related by \sqsubseteq describe session evaluations: for example, $\mu\beta?\{D \triangleright \varepsilon\}$ is *not* a possible later stage of $\mu\alpha?\{C \triangleright \alpha\}.\mu\beta?\{D \triangleright \varepsilon\}$.

We now also extend the definition of \sqsubseteq to session environments in the obvious way:

$$\Sigma \sqsubseteq \Sigma' \text{ iff } \begin{cases} \text{dom}(\Sigma) \subseteq \text{dom}(\Sigma'), \text{ and} \\ \forall k \in \text{dom}(\Sigma) : \Sigma'(k) \sqsubseteq \Sigma(k). \end{cases}$$

Below we define the judgment $h \vdash v : C$ which expresses agreement of a value with a class w.r.t. a heap, and the judgment $h \vdash \bar{v} : \tau$ which expresses agreement of a sequence of values with a session type wrt to a heap:

$$\begin{array}{c} \text{(AGGRVALUE)} \\ \hline h(v) = (C', -) \quad C' <: C \\ \hline h \vdash v : C \end{array} \quad \begin{array}{c} \text{(AGGRQUEUE-1)} \\ \hline h \vdash () : \tau \end{array} \quad \begin{array}{c} \text{(AGGRQUEUE-2)} \\ \hline h \vdash \bar{v} : \tau_i.\tau \quad h(v) = (C, -) \quad C \Downarrow \bar{C} = C_i \\ \hline h \vdash v :: \bar{v} : ?\{C \triangleright \tau\}.\tau \end{array}$$

Notice that in the judgement $\bar{v} : \tau$ the session type τ is intended as the type of a session that can safely *accept* the values in \bar{v} .

A heap is well formed iff all every live channel has its dual, and only one of them is associated to a non empty queue of values. Moreover all object fields must contain objects of the expected classes.

$$wf(h) \text{ iff } \begin{cases} k \in \text{dom}(h) \Rightarrow \tilde{k} \in \text{dom}(h), \\ h(k) \neq () \Rightarrow h(\tilde{k}) = (), \\ h \vdash o : C, fType(C, f) = C' \Rightarrow h \vdash h(o)(f) : C'. \end{cases}$$

A standard environment Γ , a session environment Σ , and a heap h agree, iff the heap is well formed, the classes of objects in the heap are subclasses of the classes associated to them by the standard environment, all subjects of the session environment occur in the heap and the session types agree with the queues of values.

$$wf(\Gamma; \Sigma; h) \quad \text{iff} \quad \begin{cases} wf(h), \\ \forall o \in \text{dom}(\Gamma): C \prec: \Gamma(o) \text{ where } h(o) = (C, -) \\ \text{dom}(\Sigma) \subseteq \text{dom}(h), \\ \forall k \in \text{dom}(\Sigma): h \vdash h(k) : \Sigma(k). \end{cases}$$

We are now able to state the Type Preservation Theorem.

Theorem 9.2 (Type Preservation). *If $wf(\Gamma; \Sigma; h)$ and $\Gamma; \Sigma \Vdash e : C$ then*

- $e, h \longrightarrow e', h'$ implies $\exists \Sigma', \Gamma'$ with $\Gamma \subseteq \Gamma'$ and $\Sigma \sqsubseteq \Sigma'$, and $wf(\Gamma'; \Sigma'; h')$, and $\Gamma'; \Sigma' \Vdash e' : C$.
- $e, h \longrightarrow e_1 \mid e_2, h'$ implies $\exists \Sigma_1, \Sigma_2, C', k, \tau, \tau'$ with $\Gamma; \Sigma_1 \Vdash e_1 : C$, and $\Gamma; \Sigma_2 \Vdash e_2 : C'$, and $wf(\Gamma; \Sigma_1; h')$, and $wf(\Gamma; \Sigma_2; h')$, and
 - either $\Sigma(k) = \tau . \tau'$ and $\Sigma_1 = \Sigma \setminus k \cup \{k : \tau'\}$ and $\Sigma_2 = \{k : \tau\}$ and $h' = h[c \mapsto ()]$ for some $c \notin h$,
 - or $\Sigma_1 = \Sigma, k : \tau$, and $\Sigma_2 = \{\tilde{k} : \tau'\}$, and $\tau \bowtie \tau'$, and $h' = h[k, \tilde{k} \mapsto ()]$ for some $k, \tilde{k} \notin h$.

The proof of this theorem can be given with a technique similar to that developed in [4].

9.2 Progress

The *progress property* assures that computations of a well typed expression cannot go wrong, *i.e.*, either a parallel of values is obtained, or a field/session of a null object is required, or a further reduction step is feasible.

In the remaining of the present section we will outline the progress proof. The present progress proof is based on those of [4, 7] but need major modifications.

Let's first give a more formal definition of progress. An expression e is *initial* if it is a user expression and $\emptyset; \varepsilon \Vdash e : C$ for some C . We will consider here only computations starting from initial expressions.

Definition 9.3. – *An expression e is a nullpointer failure if e has one of the following forms:*

- $E[\text{null}^C.f];$
- $E[\text{null}^C.s \{e\}];$
- $E[\text{null}^C.s \{k\}].$
- *An expression e has the progress property if $e, [] \longrightarrow^* P, h$ implies that*
 - *either in P all expressions are values, *i.e.*, $P = \prod_{0 \leq i < n} v_i$;*
 - *or $P, h \longrightarrow P', h'$;*
 - *or there is a subexpression of P that is a nullpointer failure.*

We call *communication* (comm for short) channels the channels which are created in rule (SessReq-R) and *synchronization* (sync for short) channels the ones created in rule (SessDel-R).

In the following we convene that the fresh communication channels created reducing a thread take successive indexes according to the order of creation, i.e. they are named $\kappa_0, \kappa_1, \dots$. This means that if $P, h \rightarrow Q, h' \rightarrow R, h''$ and κ_i is a channel created in the reduction $P, h \rightarrow Q, h'$, and κ_j is a channel created in the reduction $Q, h' \rightarrow R, h''$, then $i < j$.

Synchronization channels are created only in evaluating a delegation expression. We convene to name them with a couple of indexes: the index of the channel they send (*primary* index) and a progressive index to individuate them (*secondary* index). So c_i^0, c_i^1, \dots are the sync channels sending channels κ_i or $\tilde{\kappa}_i$. It will also be useful to decorate sync channels with the same duality as the comm channels send in the corresponding delegation session. So \tilde{c}_i^j is a sync channel created in opening a delegation session on channel $\tilde{\kappa}_i$.

The *subject* of a communication expression is the channel specified in its syntax on which the communication takes place. The index of a communication expression is the (primary) index of its subject.

- Definition 9.4.** 1. Let e be an expression and e_1, e_2 be two subexpressions of e . We say that e_1 precedes e_2 in e if, for some contexts $C[-], E[-]$ and $C'[-]$ we have $e = C[e']$ and $e' = E[e_1] = C'[e_2]$.
2. A channel ψ (where ψ can be either a comm or a sync channel) is the active channel in e if $e = E[e_1]$, ψ is the subject of a subexpression e_2 of e such that e_1 precedes e_2 in e and for no communication expressions e' we have that e_1 precedes e' and e' precedes e_2 .

Informally the active channel in an expression is the channel on which the next communication will be performed. Obviously if e is a communication expression the active channel of $E[e]$ is the subject of e itself. We say that a subexpression e' is the last subexpression in e if e' does not precede any other subexpression of e .

Lemma 9.5. Let $e, [] \rightarrow^* P, h$ where e is an initial expression. Then

1. No thread in P can contain occurrences of both κ and $\tilde{\kappa}$ for some channel κ .
2. A comm channel can be active in only one thread of P .
3. A sending expression with a sync subject c_i^j can occur only as last expression of a thread in P which was opened by a delegation expression on κ_i .

Proof. By induction on reduction.

Lemma 9.6. Let $e, [] \rightarrow^* e' \mid P, h$ where e is an initial expression. Then the index of the active channel in e' is greater than or equal to the (primary) index of any other channel occurring in e' .

Proof. Induction on reduction. Note that when a new thread is opened by delegation the sync channel takes the same index of the sent channel.

Theorem 9.7 (Progress). *If e is an initial expression, then e has the progress property.*

Proof. If e is initial we have $\emptyset; \varepsilon \vdash e : C$. Assume that e does not have the progress property. Then $e, [] \longrightarrow^* e_1 \mid \dots \mid e_n, h$ which is irreducible and is not a parallel composition of values and it does not contain nullpointer failures. By the subject reduction property we have that there are $\Gamma, \Sigma_1, \dots, \Sigma_n$ and C_1, \dots, C_n such that $wf(\Gamma; \Sigma_i; h)$ and $\Gamma; \Sigma_i \Vdash e_i : C_i$ for all $1 \leq i \leq n$. We can assume without loss of generality that e_1, \dots, e_m are not values. The evaluation can only be stopped by a receiving expression waiting for data in the associated channel. So for all $1 \leq i \leq m$ we must have $e_i = E[e'_i]$ where e'_i is a receiving expression whose subject is the active channel in e_i . Let's say that e_i is *blocked* on the subject of e'_i . First note that in e_1, \dots, e_m the following property hold:

(P) If an expression e_i is blocked on a sync channel c_p^l then there must exist another expression e_j which is blocked on a comm channel k_q with $q \geq p$.

Proof. Assume that e_i is blocked on a sync channel c_p^l . Then the (unique) sending expression on c_p^l must occur in some other expression e_j which, by Lemma 9.6, is blocked on a communication expression whose index is greater than or equal to p . If its subject is a comm channel we are done. Otherwise an iterative argument concludes the proof.

Let's consider an expression e'_i such that its subject is a channel k_p where p is the highest among the index of channels which are subjects of e'_1, \dots, e'_m . Property **(P)** assures that such an expression must exist. Since the corresponding queue in the heap is empty and all expressions are well typed and consistent with the heap, then there must exist a corresponding sending expression e' whose subject is \tilde{k}_p and which occur in e_j for some $j \neq i$ by Lemma 9.5.1. Now we have two cases:

1. e_j is blocked on a receiving comm expression on \tilde{k}_p (by definition p is the highest index of blocked channels);
2. e_j is blocked on a receiving expression on a sync channel \tilde{c}_p^q .

Case 1 is impossible by the well typing and well forming heap conditions. In case 2, by Lemma 9.5.3 we have that there must exist a sending expression with subject \tilde{c}_p^q which is the last expression of a blocked expression e_l opened by a delegation on \tilde{k}_p . Since p is the greatest index, then the active channel of e_l must be either \tilde{k}_p or a sync channel \tilde{c}_p^l . In the former case we conclude as in case 1, in the latter case we conclude by an iterative argument.

10 Summary, Conclusions, and Further Work

We believe that a good combination of session types with object oriented programming will be a crucial vehicle for the development of safe concurrent and/or distributed, collaborating programs. We believe that such a combination should be achieved through *amalgamation* rather than extension of the paradigms.

We have suggested such an amalgamation, whereby sessions subsume methods, threads consist of the execution of session bodies on objects and communicate with each other through asynchronously sending/receiving objects on channels, where all choices are based on the classes of objects, and where data are exchanged between

threads running sessions via channels using asynchronous communication, and where sessions can be delegated to other sessions.

One major deviation of ST00P wrt other sessions languages, is that ST00P sessions are not first class, and instead we are using the concept of delegation. This significantly simplifies the system, although it restricts the expressive power. We believe that this restriction has small practical relevance, *i.e.*, that most useful applications can be written with this restriction.

We want to investigate the expressive power of the paradigm in terms of more examples. We also want to explore how slightly extend the expressivity of delegation, so that it supports an initial and a final dialogue before and after the delegation. In terms of the example from Section 3, the seller might have an initial dialogue with the bank before delegating *e.g.*, giving precise instruction for the credit check, and might have a final dialogue after the credit check, *e.g.*, receiving more detailed information about the mode of payment.

After that, we plan to explore full language designs based on ST00P, *e.g.*, adding concurrency and synchronization for kernel applications, or, alternatively considering distribution and data duplication issues for web services applications.

Acknowledgements

We are greatly indebted to Dimitris Mostrous for many interesting discussions which, among others, lead to the idea of amalgamating sessions with methods. We are very grateful to Nobuko Yoshida for extensive feedback.

References

1. Eduardo Bonelli, Adriana Compagnoni, and Elsa Gunter. Correspondence Assertions for Process Synchronization in Concurrent Communications. *Journal of Functional Programming*, 15(2):219–248, 2005.
2. Marco Carbone, Kohei Honda, and Nobuko Yoshida. A Theoretical Basis of Communication-centered Concurrent Programming. Web Services Choreography Working Group mailing list, to appear as a WS-CDL working report.
3. Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured Communication-Centred Programming for Web Services. In Rocco De Nicola, editor, *ESOP'07*, volume 4421 of *LNCS*, pages 2–17. Springer-Verlag, 2007.
4. Mario Coppo, Mariangiola Dezani-Ciancaglini, and Nobuko Yoshida. Asynchronous Session Types and Progress for Object-Oriented Languages. In Marcello Bonsangue and Einar Broch Johnsen, editors, *FMOODS'07*, volume 4468 of *LNCS*, pages 1–31. Springer-Verlag, 2007.
5. Ferruccio Damiani, Elena Giachino, Paola Giannini, Nick Cameron, and Sophia Drossopoulou. A State Abstraction for Coordination in Java-like Programming. *Acta Informatica*, 2007. To appear.
6. Mariangiola Dezani-Ciancaglini, Sophia Drossopoulou, Elena Giachino, and Nobuko Yoshida. Bounded Session Types for Object-Oriented Languages. In Frank de Boer and Marcello Bonsangue, editors, *FMCO'06*, volume 4709 of *LNCS*, pages 207–245. Springer-Verlag, 2007.

7. Mariangiola Dezani-Ciancaglini, Dimitris Mostrous, Nobuko Yoshida, and Sophia Drossopoulou. Session Types for Object-Oriented Languages. In Dave Thomas, editor, *ECOOP'06*, volume 4067 of *LNCS*, pages 328–352. Springer-Verlag, 2006.
8. Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, Alex Ahern, and Sophia Drossopoulou. A Distributed Object Oriented Language with Session Types. In Rocco De Nicola and Davide Sangiorgi, editors, *TGC'05*, volume 3705 of *LNCS*, pages 299–318. Springer-Verlag, 2005.
9. Manuel Fähndrich, Mark Aiken, Chris Hawblitzel, Orion Hodson, Galen C. Hunt, James R. Larus, , and Steven Levi. Language Support for Fast and Reliable Message-based Communication in Singularity OS. In Willy Zwaenepoel, editor, *EuroSys2006*, ACM SIGOPS, pages 177–190. ACM Press, 2006.
10. Pablo Garralda, Adriana Compagnoni, and Mariangiola Dezani-Ciancaglini. BASS: Boxed Ambients with Safe Sessions. In Michael Maher, editor, *PPDP'06*, pages 61–72. ACM Press, 2006.
11. Simon Gay and Malcolm Hole. Subtyping for Session Types in the Pi-Calculus. *Acta Informatica*, 42(2/3):191–225, 2005.
12. Simon Gay, Vasco T. Vasconcelos, and António Ravara. Session Types for Inter-Process Communication. TR 2003–133, Department of Computing, University of Glasgow, 2003.
13. Kohei Honda. Types for Dyadic Interaction. In Eike Best, editor, *CONCUR'93*, volume 715 of *LNCS*, pages 509–523. Springer-Verlag, 1993.
14. Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language Primitives and Type Disciplines for Structured Communication-based Programming. In Chris Hankin, editor, *ESOP'98*, volume 1381 of *LNCS*, pages 22–138. Springer-Verlag, 1998.
15. Kohei Honda, Nobuko Yoshida, and Marco Carbone. Web Services, Mobile Processes and Types. *EATCS Bulletin*, 91:160–188, 2007.
16. Dimitris Mostrous and Nobuko Yoshida. Two Sessions Typing Systems for Higher-Order Mobile Processes. In Simona Ronchi Della Rocca, editor, *TLCA'07*, volume 4583 of *LNCS*, pages 321–335. Springer-Verlag, 2007.
17. Stephen Sparkes. Conversation with Steve Ross-Talbot. *ACM Queue*, 4(2):14–23, 2006.
18. Kaku Takeuchi, Kohei Honda, and Makoto Kubo. An Interaction-based Language and its Typing System. In C. Halatsis, D. Maritsas, G. Philokyprou, and S. Theodoridis, editors, *PARLE'94*, volume 817 of *LNCS*, pages 398–413. Springer-Verlag, 1994.
19. Antonio Vallecillo, Vasco T. Vasconcelos, and António Ravara. Typing the Behavior of Objects and Components using Session Types. In Antonio Brogi and Jean-Marie Jacquet, editors, *FOCLASA'02*, volume 68(3) of *ENTCS*, pages 439–456. Elsevier, 2002.
20. Vasco T. Vasconcelos, Simon Gay, and António Ravara. Typechecking a Multithreaded Functional Language with Session Types. *Theoretical Computer Science*, 368(1-2):64–87, 2006.
21. Web Services Choreography Working Group. Web Services Choreography Description Language. <http://www.w3.org/2002/ws/chor/>.