



FLORE

Repository istituzionale dell'Università degli Studi di Firenze

Session and Union Types for Object Oriented Programming

Questa è la Versione finale referata (Post print/Accepted manuscript) della seguente pubblicazione:

Original Citation:

Session and Union Types for Object Oriented Programming / L. Bettini; S. Capecchi; M. Dezani-Ciancaglini; E. Giachino; B. Venneri.. - STAMPA. - (2008), pp. 659-680. [10.1007/978-3-540-68679-8_41]

Availability:

This version is available at: 2158/319871 since: 2017-05-26T11:19:46Z

Publisher: Springer

Published version: DOI: 10.1007/978-3-540-68679-8_41

Terms of use: Open Access

La pubblicazione è resa disponibile sotto le norme e i termini della licenza di deposito, secondo quanto stabilito dalla Policy per l'accesso aperto dell'Università degli Studi di Firenze (https://www.sba.unifi.it/upload/policy-oa-2016-1.pdf)

Publisher copyright claim:

(Article begins on next page)

Session and Union Types for Object Oriented Programming*

Lorenzo Bettini¹, Sara Capecchi¹, Mariangiola Dezani-Ciancaglini¹, Elena Giachino¹, and Betti Venneri²

¹ Dipartimento di Informatica, Università di Torino
 ² Dipartimento di Sistemi e Informatica, Università di Firenze

Dedicated to Ugo Montanari on the Occasion of his 65th Birthday

Abstract. In network applications it is crucial to have a mechanism to guarantee that communications evolve correctly according to the agreed protocol. Session types offer a method for abstracting and validating structured communication sequences (sessions). In this paper we propose union types for refining and enhancing the flexibility of session types in the context of communication centred and object oriented programming. We demonstrate our ideas through an example and a calculus formalising the main issues of the present approach. The type system garantees that, in well-typed executable programs, after a session has started, the values sent and received will be of the appropriate type, and no process can get stuck forever.

Keywords: Sessions, Object Oriented Programming, Session Types, Union Types.

1 Introduction

Writing safe communication protocols has become a central issue in the theory and practice of concurrent and distributed computing. The actual standards still leave to the programmer much of the responsibility in guaranteeing that the communication will evolve as agreed by all the involved agents.

Session types [26, 27] offer a method for abstracting and validating structured communication sequences (sessions). This is achieved 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. Since the specification of a session is a type, the conformance test of programs with respect to specifications becomes type checking.

The popularity of class-based *object oriented* languages justifies the interest in searching for class definitions which naturally include communication primitives. For

^{*} This work has been partially supported by MIUR project EOS DUE and by EU Project Software Engineering for Service-Oriented Overlay Computers (SENSORIA, contract IST-3-016004-IP-09).

P. Degano et al. (Eds.): Montanari Festschrift, LNCS 5065, pp. 659-680, 2008.

[©] Springer-Verlag Berlin Heidelberg 2008

this reason, an amalgamation of communication centred and object oriented programming has been first proposed in [18], where methods are unified with sessions and choices are based on the classes of exchanged objects.

Union types have been shown useful for enhancing the flexibility of subtyping in various settings [1, 21, 11, 10, 29]. For example a bank can answer yes or not according to the balance between an account and an item price. If yes and not are objects of classes OK and NoMoney respectively, then the class of the object answer is naturally the union of the two classes OK and NoMoney, *i.e.* OK \lor NoMoney. Without union types typing answer would require a superclass of both OK and NoMoney to be already defined, and this superclass could include unwanted objects. With union types we can express communications between parties which manipulate heterogeneous objects just by sending and receiving objects which belong to subclasses of one of the classes in the union. In this way the flexibility of object-oriented depth-subtyping is enhanced, by strongly improving the expressiveness of choices based on the classes of sent/received objects.

The aim of the present paper is to discuss and formalise the use of union types for session-centred communications in a core object-oriented calculus. A preliminary version of the basic calculus, without union types, is defined in [18]. In the present paper, the calculus of [18] is formally revised, so that typing and semantics are rather cleaner and simpler. Furthermore, the extension to union types, which is the main novelty of the present proposal, poses specific problems in formulating reduction and typing rules to ensure that communications are safe while flexible.

We first present an example which illustrates the main features of our approach and then we formalise these features through a featherweight representation. We call SAM^{\vee} (Sessions Amalgamated with *M*ethods plus union types) the language of the example and $\mathscr{F}SAM^{\vee}$ the formalising calculus.

SAM^{\vee} Overview. SAM^{\vee}, as the language of [18], is concerned with the amalgamation of the object oriented features with the session part, but it is agnostic w.r.t. to the remaining features of the language, such as whether the language is distributed or concurrent, and the features for synchronisation.

In SAM^{\lor}, sessions and methods are "amalgamated": invocation takes place on an object and the execution takes place immediately and concurrently with the requesting thread (indeed, SAM^{\lor} is multi-threaded and the communication is asynchronous). Thus, it keeps the method-like invocation mechanism while involving two threads, typical of session based communication mechanisms. The body is determined by the class of the receiving object (avoiding in this way the usual branch/select primitives [27]), and any number of communications interleaved with computation is possible. Sessions are defined in a class, which can have also fields. We believe that the above amalgamated model of session naturally reflects our intuition of services. Furthermore, it can neatly encode "standard" methods.

A thread can make a session request through e.s $\{e'\}$, where e is an expression denoting an object, s is the name of a session defined in the object's class; then, e is evaluated to an object o, and the session body of s in o's 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'^1 . Notice that channels are implicit, and are not written by the programmer. At every step, in each thread, there is only one single active channel on which communications are performed.

The expressions send(e) and rec(x) send and receive objects on the active channel, respectively. The expression $sendC(e)\{C_1 \Rightarrow e_1 \parallel ... \parallel C_n \Rightarrow e_n\}$ (where C means Case) evaluates e to an object and sends it on the active channel, and then continues with e_i , where C_i is the class that best fits the class of the object sent. The counter part of sendCis the expression $recC(x)\{C_1 \Rightarrow e_1 \parallel ... \parallel C_n \Rightarrow e_n\}$, where the choice is based on the class of the object received. The expression $sendW(e)\{C_1 \Rightarrow e_1 \parallel ... \parallel C_n \Rightarrow e_n\}$ (where W means While) is similar to $sendC(e)\{C_1 \Rightarrow e_1 \parallel ... \parallel C_n \Rightarrow e_n\}$, except that it allows for enclosed cont, which continues the execution at the nearest enclosing sendW. The expression $recW(x)\{C_1 \Rightarrow e_1 \parallel ... \parallel C_n \Rightarrow e_n\}$ has the obvious meaning. Finally, $e \bullet s$ {} delegates the current session to the object resulting from the evaluation of e; 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.

Related Papers. We describe $\mathscr{F}SAM^{\vee}$ following *Featherweight Java* [30], which today has become a standard for class based object calculi.

Session Types have been first introduced to model communication protocols between π -calculus processes [26, 32, 27]. They have been made more expressive by enriching them with correspondence assertions [3], subtyping [24], bounded polymorphism [23] and safer by assuring deadlock-freedom [14]. More recently session types have been extended to multi-party communications [2, 9].

Session types have been developed also for CORBA [33], for functional languages [25, 34], for boxed ambients [22], for the W3C standard description language for Web Services CDL [8, 35, 31, 28], for operating systems [19], and for object oriented programming languages [17, 16, 13, 15, 18, 6]. In [6] generic types are added to a language/calculus based on the approach of [18]; independently from the different typing extensions, $\mathscr{F}SAM^{\vee}$ improves the definition of the calculi of [18] and [6], both in syntax and in operational semantics.

Union types have been proved useful for functional languages [1, 11], for objectoriented languages [29], for languages manipulating semi-structured data [21] and for the π -calculus [10]. We will tell more on the relations between the present paper and [29] at the end of Subsection 6.1.

There are many concurrent object-oriented languages and calculi in the literature; for this topic we refer to the related work section of [20].

Paper Structure. In Section 2 we describe SAM^{\vee} in terms of an example. We then proceed by formalising the calculus $\mathscr{F}SAM^{\vee}$, its typing and semantics. Section 7 draws some future work directions.

2 An Example

In this section, we describe SAM^{\vee} through an example, which expresses a typical collaboration pattern, *c.f.* [35, 7, 8], and which refines the example of [18]. This simple

 $^{^{-1}}$ k and \tilde{k} play for channels a role similar to that of this for objects.

```
sessiontype Shopping_ST =
1
       !Item.?Money.\mu\alpha.!{ OK \Rightarrow !AccntNr.!Money.?{OK \Rightarrow ?Date, NoMoney \Rightarrow \varepsilon },
2
                                    NoDeal \Rightarrow arepsilon ,
3
                                    MakeAnOffer \Rightarrow ? { Money \Rightarrow \alpha, NoDeal \Rightarrow \varepsilon } }
4
5
     sessiontype ExaminePrice_ST = ?Money.!Ok\NoDeal\MakeAnOffer
6
7
     sessiontype Sell_ST =
8
      ?Item.!Money.\mu\alpha.?{ OK \Rightarrow ?AccntNr.?Money.!{ OK\Rightarrow!Date, NoMoney \Rightarrow \varepsilon },
9
                                    NoDeal \Rightarrow \epsilon,
10
                                    MakeAnOffer \Rightarrow ! { Money \Rightarrow \alpha, NoDeal \Rightarrow \varepsilon } }
11
12
     sessiontype CalDelDate_ST = ?Item.!Date
13
14
     sessiontype CalNewPrice_ST = ?Money.!Money\NoDeal
15
16
     sessiontype CreditCheck_ST = ?AccntNr.?Money
17
```

Fig. 1. Session types for the buyer-seller example

protocol contains essential features which demonstrate the expressiveness of the idioms of SAM^{\vee} . 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, SAM^{\vee} can delegate the current session through a session call as in [18,6]. The negotiation allows several rounds: the buyer may either accept the price, or break the negotiation, or require a better deal by sending different kinds of answers; in the latter case, the seller might respond by sending a better price, or might break the negotiation sending a negative answer. Thus, branch selection in control structures is based on the dynamic class of an object sent.

The session types Shopping_ST and Sell_ST (see Figure 1) describe the communication pattern between the Buyer and the Seller.

The session type Shopping_ST 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 $\mu \alpha$.! { $OK \Rightarrow ..., NoDeal \Rightarrow ..., MakeOnOffer \Rightarrow ... }$ 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 negotiation resumes on the basis of it, whereas if it is a NoDeal, the negotiation ends.

Note that in both Shopping_ST and Sell_ST the recursion variable is nested inside multiple choices, so that this behaviour could not have been expressed using regular expressions as in [15]. The use of recursive types has also other advantages, like that of allowing iterative expressions with multiple exit points and multiple recursions. In

```
class Buyer {
    AccntNr accnt; Seller seller;
2
3
    String VDate Shopping_ST shopping
4
         { Item prodId := ....;
5
            seller.sell{
6
               send(prodID);
7
               Money price := rec;
8
               examinePrice{send(price);
9
                              OK\VDeal\MakeAnOffer resp := rec};
10
               sendW(resp){
11
                  OK \Rightarrow \{ send(accnt); send(price); \}
12
                          recC(x) { OK \Rightarrow Date delivDate:=rec;
13
                                          NoMoney \Rightarrow new String("no money"); } }
14
                  NoDeal \Rightarrow new String("refusing proposed price")
15
                  MakeAnOffer \Rightarrow {
16
                           recC(x) { Money \Rightarrow examinePrice{send(x);
17
                                                          resp := rec};
18
                                                          cont;
19
                                            NoDeal \Rightarrow new String("offer refused")}
20
                                      }
21
                          } //end of session call sell
22
                      } //end of session shopping
23
24
   Object ExaminePrice_ST examinePrice
25
         { Money price := rec
26
          ... //code for decision
27
          send(resp)
28
         } // end of session examinePrice
29
   3
30
```

Fig. 2. The class Buyer

Figure 2 we show the implementation of the class Buyer. It has the fields accnt and seller, which will contain the account number and the seller used to buy products.

The class Buyer supports two sessions called shopping and examinePrice. Session shopping has session type Shopping_ST and return type String\Date. The union type String\Date describes the possible results of the negotiation: in case of success the session ends returning the date of the delivery of the item; in case of failure it returns a string describing the reason. In the body of this session the desired product is determined and stored in prodId (line 5). Then, a session request is made to the seller to run session sell (line 6). 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 7), receive a Money and store it in price (line 8). Based on its value the buyer will calculate his response calling session examinePrice which returns the answer in resp (lines 9 and 10). Let us notice that resp's type is the union of all the possible answers' type. On line 11, the buyer enters a loop with sendW, where he sends resp, and branches according to its class. If resp is OK, indicating acceptance of the price,

```
class Seller {
1
      Bank bank;
2
3
      String VItem Sell_ST sell
4
          { Item prodID := rec;
5
           Money price=...;
6
           send(price);
7
           recW( x ){
8
                 OK \Rightarrow \{ sendC( bank \cdot check \{ \} ) \}
9
                               \{OK \Rightarrow calDelDate\{send(prodID);
10
                                                    Date date := rec};
11
12
                                       send(date); prodID []
                                NoMoney \Rightarrow new String("failed bank transaction");
13
                                      } } [
                 NoDeal \Rightarrow new String("proposal refused by buyer")
14
                 MakeAnOffer \Rightarrow { calNewPrice{send(price);
15
                                        Money \lor NoDeal resp := rec};
16
                                        sendC(resp) { Money \Rightarrow cont \parallel
17
                                                     NoDeal \Rightarrow new String("refusing
18
                                                          proposed price"); } }
             }
19
         } //end of session sell
20
21
      Object CalDelDate_ST calDelDate
22
        { Item item := rec
23
           ... //code for calculate date
24
          send(date)
25
         } // end of session calDelDate
26
27
      Object CalNewPrice_ST calNewPrice
28
29
        { Money price := rec
           ... //code for response (if the answer is positive then the field
30
             //price is updated with the new price)
31
          sendC(resp)
32
          } // end of session calNewPrice
33
      }
34
```

Fig. 3. The class Seller

then the buyer will send his account 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). In this case the reason of failure is stored in the string failure. 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, indeed Line 17 contains recC 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 session examinePrice is called, which returns the buyer's reaction in resp, and the recursion will continue (line 19). If the value received is NoDeal, then the loop will be abandoned.

```
class Bank {
1
       Ok VNoMoney CreditCheck_ST check
2
          { AccntNr accnt := rec;
3
             Money amt := rec;
4
             .... // code for check
5
             If(response) then new Ok else new NoMoney;
6
          }// end of session check
7
     }
8
```

Fig. 4. The class Bank

Notice that in order to get an arbitrary number of repetitions, it is crucial to allow objects of different classes to be sent in the different iterations of while loops.

The session type Sell_ST describes the protocol from the point of view of the seller, and is "dual" to Shopping_ST. We now consider the class Seller, from Figure 3. 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 recW loop, which is the counterpart to the sendW loop from shopping and performs all the seller's negotiation. The interesting 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. At the end of the execution of check the session will continue according to the bank answer (OK or NoMoney).

The session type for check from class Bank in Figure 4 is the receipt of a AccntNr and a Money followed by sending either OK, or a NoMoney object. Note that the 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 union of the types of the possible answers, *i.e.* OK \vee NoMoney.

Notice that the sessions examinePrice, calNewDate and calNewPrice are examples of the implementation in SAM^V of methods, since they start by receiving arguments and after elaborating them send a result.

3 Syntax

This section presents the syntax of $\mathscr{F}SAM^{\vee}$ (Figure 5), a minimal concurrent and imperative core calculus, based on *Featherweight Java* [30] (abbreviated with FJ). $\mathscr{F}SAM^{\vee}$ supports the basic object-oriented features and session request, session delegation, branching sending/receiving and loops. In details, $\mathscr{F}SAM^{\vee}$ encompasses the following linguistic features: basic object oriented expressions, session bodies and communication constructs that combine send/receive with branching and loops.

We use grey to indicate expressions that are produced during the reduction process, but do not occur in the source code of a program. We also use the standard convention of denoting with $\overline{\xi}$ a sequence of elements $\xi_1, ..., \xi_n$.

Union types are defined as in [29]: they are built out of class names by the union operator (denoted by \vee).

Programs are defined from a collection of classes. The metavariables C and D, possibly with subscripts, range over class names. Each class has a name, a list of *fields* of the

```
(union type)
                              ::=
                                       C \mid T \lor T
                           Т
(class)
                                       class C \triangleleft C \{ \overline{Tf}; \overline{S} \}
                           L
                               ::=
(session)
                           S ::=
                                       Tts{e}
(expression)
                                       x | this | cont | o | e; e | e.f:= e | e.f | new C(\overline{e})
                           e
                               ::=
                                        e.s{e} | e \bullet s{k}
                                        k. sendC(e){C \Rightarrow e || C \Rightarrow e}
                                         k. recC(x) \{ C \Rightarrow e \mid | C \Rightarrow e \}
                                         k. sendW(e){C \Rightarrow e || C \Rightarrow e}
                                         k. recW(x){C \Rightarrow e || C \Rightarrow e}
(parallel threads) P
                                         e | P || P
                                ::=
```

Fig. 5. Syntax, where syntax occurring only at runtime appears shaded. Syntax for session types t is in Figure 9.

form \overline{Tf} , where f represents the field name and T its type, and a list of *sessions* of the form $Tts \{e\}$, where T is the return type, t the session type, s the session name, and e the session body. For the sake of conciseness the symbol \triangleleft represents class extension, as in [30]. All classes are defined as extensions of the topmost class Object.

Expressions include variables, that are both standard term variables x and the special variables this and cont. The variable this is considered implicitly bound in any session declaration. Instead, sendW and recW are the only binders for cont, that represents the continuation by recursive computation. Let us notice that free occurrences of cont in e are not bound in the expression sendW(e){...}: actually no occurrence of cont can appear in e if this expression is typable (see rule SENDW-T in Figure 11).

In a *session request* $e.s \{e'\}$ we call the expression e' the *co-body* of the request (since it will be evaluated concurrently with the body of requested session).

In the *session delegation* expression, $e \bullet s \{k\}$, the channel k is added by the operational semantics in order to keep track of the channel to pass to the delegated session.

Channels are implicit in the source language syntax. At runtime, communication channels k are introduced at each new session request. We denote the dual with ..., where \tilde{k} is again a runtime channel, and where ... is an involution: $\tilde{k} = k$. Whenever a thread uses a channel k, the other participant in the communication uses its dual \tilde{k} . The operational semantics associates to k and \tilde{k} two different queues of messages; when a thread, which uses the channel k, wants to receive a message it will inspect the queue associated to k, while, when it sends a message it will add it to the queue associated to \tilde{k} (see Section 5).

The body of a *communication expression* is a pair of alternatives $\{C_1 \Rightarrow e_1 \mid C_2 \Rightarrow e_2\}$, whose choice depends on the class of the object that is sent or received. In particular, in case of sendW and recW the expressions e_i can contain cont, representing recursive computations.

With respect to SAM^{\lor} we only have binary choices as bodies of communication expressions in $\mathscr{F}SAM^{\lor}$, since the other forms can be encoded with them. First of all, a unary choice {C \Rightarrow e} can be simply encoded as {C \Rightarrow e [] C \Rightarrow new Object()}, since sending or receiving an object of class C will always choose the first alternative. With

unary choices we can encode the two communication constructs used in the example, i.e., send for sending, and rec for receiving, that we omit from $\mathscr{F}SAM^{\vee}$: the expression sendC(e){Object \Rightarrow new Object()} encodes send(e) and in a similar way, recC(x){Object \Rightarrow x} encodes rec.

With binary choices we can also encode *n*-ary choices for n > 2, getting in this way the constructs used in the example of Section 2. The informal idea of such encoding, given the set of classes to be used in the choices, is to take the first (in the left to right order) relative minimum from this set and use it to define the first branch; the second branch will use Object and we iterate this procedure to write the expression of this second branch, until we remain with only two choices. Thus, for instance, consider the choice $\{C_1 \Rightarrow e_1 || C_2 \Rightarrow e_2 || C_3 \Rightarrow e_3\}$, where C_1 is not related to the other classes and C_2 is a subclass of C_3 ; we can encode this choice with the (nested) binary choice: $\{C_1 \Rightarrow e_1 || Object \Rightarrow \{C_2 \Rightarrow e_2 || C_3 \Rightarrow e_3\}$. Notice that this encoding is correct also if $C_2 = C_3$.

The types used for selecting branches in a choice are class names. This simplifies the formal treatment and the proofs, but, again, we can encode choices with arbitrary union types by *n*-ary choices in a straightforward way; e.g., $\{C_1 \lor C_2 \Rightarrow e \mid C_3 \Rightarrow e'\}$ is encoded as $\{C_1 \Rightarrow e \mid C_2 \Rightarrow e \mid C_3 \Rightarrow e'\}$.

A *runtime expression* is either a user expression (i.e. an expression in Figure 5 without shaded syntax) or an expression containing channels and/or object indentifiers. Furthermore, threads of runtime expressions can occur at runtime (see the operational semantics). Parallel threads are ranged over by P. Fully evaluated objects will be represented by object identifiers denoted by o.

The main novelty of $\mathscr{F}SAM^{\vee}$ w.r.t. FJ is that session invocation can involve the creation of concurrent and communicating threads. Other minor differences are: we do not have cast and overriding, which are orthogonal to our approach; we do not have explicit constructors, then in the object instantiation expression new $C(\overline{e})$, the values \overline{o} to which \overline{e} reduce are the initial values of the fields.

Notice that standard methods can be seen as special cases of sessions. In fact, a method declaration can be (informally) encoded as a session with nested recCs (one for each parameter) and with one sendC returning the method body. Similarly, method calls are special cases of session requests: the passing of arguments is encoded as nested sendCs (one for each argument) and the object returned by the method body is retrieved via one recC. This encoding will use unary choices that can be rendered with binary choices as explained above.

4 Auxiliary Functions

As in FJ, a class table *CT* is a mapping from class names to class declarations with domain $\mathscr{D}(CT)$. Then a program is a pair (CT, e) of a class table (containing all the class definitions of the program) and an expression e (an expression belonging to the source language representing the program's main entry point). The class Object has no fields and its declaration does not appear in *CT*. As in FJ, from any *CT* we can read off the subtype relation between classes, as the transitive closure of \triangleleft clause; moreover this relation is extended in order to relate types built out of union (Figure 6). As usual considering union types modulo the equivalence relation induced by <: we get the commutativity and associativity of \lor . Therefore each union type can be written

$$\begin{array}{c|c} T <: T & \frac{T <: T' & T' <: T''}{T <: T''} & \frac{c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\} \in CT}{C <: D} \\ \hline T <: T \lor T' & T' <: T \lor T' & \frac{T' <: T & T'' <: T}{T' \lor T'' <: T} \\ fields(Object) = \bullet & \frac{fields(D) = \overline{T'f'} & c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\}}{fields(C) = \overline{Tf}} \\ fields(Object) = \bullet & \frac{fields(D) = \overline{T'f'} & c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\}}{fields(C) = \overline{Tf}} \\ figure_w(f_i, C) = ftype_r(f_i, C) = T_i \\ ftype_w(f_i, C) = ftype_r(f_i, C) = T_i \\ ftype_w(f, T_1 \lor T_2) = \begin{cases} ftype_w(f, T_1) & \text{if } ftype_w(f, T_2) <: ftype_w(f, T_2), \\ ftype_w(f, T_1 \lor T_2) = ftype_r(f, T_1) \lor ftype_r(f, T_2) \\ (\bot & \text{otherwise}. \\ ftype_r(f, T_1 \lor T_2) = ftype_r(f, T_1) \lor ftype_r(f, T_2) \\ c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\} & \operatorname{Tts} \left\{ e \right\} \in \overline{S} \\ stype(s, C) = \left\{ t \right\} & c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\} & s \notin \overline{s} \\ rtype(s, C) = T & rtype(s, C) \\ rtype(s, T_1 \lor T_2) = rtype(s, T_1) \lor type(s, T_2) \\ rtype(s, C) = rtype(s, T_1) \lor rtype(s, T_2) \\ c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\} & \operatorname{Tts} \left\{ e \right\} \in \overline{S} \\ rtype(s, C) = rtype(s, T_1) \lor rtype(s, T_2) \\ rtype(s, C) = rtype(s, T_1) \lor rtype(s, T_2) \\ rtype(s, C) = e & c \operatorname{lass} C \lhd D \left\{ \overline{Tf}; \overline{S} \right\} & s \notin \overline{S} \\ sbody(s, C) = sbody(s, D) \end{cases}$$

Fig. 6. Subtyping and Lookup Functions

as $C_1 \vee \ldots \vee C_n$ for $n \ge 1$: we say that the classes C_1, \ldots, C_n build the union type $C_1 \vee \ldots \vee C_n$. A union type $C_1 \vee \ldots \vee C_n$ is proper if n > 1.

We assume a fixed *CT* that satisfies some usual sanity conditions as in FJ [30]. Thus, in the following, instead of writing $CT(C) = class \dots$ we will simply write class C

We define auxiliary functions (see Figure 6) to lookup fields and sessions from CT; these functions are used in the typing rules and in the operational semantics. The *fields* lookup function is as in FJ. As for field type lookup we distinguish between the contexts where the field is used for reading (*ftype_r*) from those where it is for writing (*ftype_w*). The *stype* and *rtype* return a set of session types and the return type of a session, respectively, while *sbody* returns the body of a session. As in FJ these functions may have to inspect the class hierarchy in case the required element is not present in the current class.

Notice that the type lookup functions take a type as argument (not simply a class name) because the receiver expression of a field/session access may be of a proper union type. As for field type lookup, when the field is used in read mode, in case of a proper union type, we simply return the union type of the result of $ftype_r$ invoked on the argument types (if both retrievals succeed). On the contrary, when a field is updated, due to the contravariance relation, in case of a proper union type we must

$$e\langle k\rangle = \begin{cases} e_1\langle k\rangle; e_2\langle k\rangle & \text{if } e = e_1; e_2, \\ e_1\langle k\rangle.f & \text{if } e = e_1.f, \\ e_1\langle k\rangle.f:=e_2\langle k\rangle & \text{if } e = e_1.f:=e_2, \\ e_1\langle k\rangle.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\langle k\rangle\bullet s\{k\} & \text{if } e = e_1\bullet s\{\}, \\ k.\text{sendC}(e_0\langle k\rangle)\{\overline{C} \Rightarrow e\langle k\rangle\} & \text{if } e = \text{sendC}(e_0)\{\overline{C} \Rightarrow e\}, \\ k.\text{sendC}(x)\{\overline{C} \Rightarrow e\langle k\rangle\} & \text{if } e = \text{sendC}(x)\{\overline{C} \Rightarrow e\}, \\ k.\text{sendW}(e)\{\overline{C} \Rightarrow e\langle k\rangle\} & \text{if } e = \text{sendW}(e)\{\overline{C} \Rightarrow e\}, \\ k.\text{sendW}(e)\{\overline{C} \Rightarrow e\langle k\rangle\} & \text{if } e = \text{sendW}(e)\{\overline{C} \Rightarrow e\}, \\ k.\text{recC}(x)\{\overline{C} \Rightarrow e\langle k\rangle\} & \text{if } e = \text{recW}(x)\{\overline{C} \Rightarrow e\}, \\ e & \text{otherwise.} \end{cases}$$

$$e\lfloor e'/\text{cont} \rfloor = \begin{cases} e_1\lfloor e'/\text{cont} \rfloor; e_2\lfloor e'/\text{cont} \rfloor & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.f:=e_2\lfloor e'/\text{cont} \rfloor & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2\} & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2] & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2] & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2] & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2] & \text{if } e = e_1.s\{e_2\}, \\ e_1\lfloor e'/\text{cont} \rfloor.s\{e_2] & \text{if } e = e_1.s\{e_2\}, \\ e_1 = e_1.s\{e_2\},$$

Fig. 7. Channel Addition and Continuation Replacement

return the intersection of the result of $ftype_w$ on the arguments; however, in the absence of multiple inheritance, the only possible cases are those listed in Figure 6, thus we can avoid introducing intersection types.

As for the *stype* lookup function, it returns a set of session types; in case it is invoked with a class name as argument, it will return a singleton. The interesting case is when it is invoked with a proper union type: it will return the union of the sets corresponding to the argument types, so that we have all the session types of the classes that build the union type (see how it is used in the typing system, Figures 11 and 13). The *rtype* lookup function behaves in a covariant way since the resulting object cannot be used in writing mode. We notice that *sbody* is only invoked with a class name as type argument, since we invoke sessions on objects only, and all objects have a class name as type.

It is easy to verify that all lookup functions applied to equivalent union types return either equivalent union types or the same sets of session types, whenever they are defined.

5 Operational Semantics

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

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

where :: denotes heap concatenation.

During evaluation, any expression new $C(\overline{o})$ will be replaced by a new object identifier o. The heap will then maps the object identifier o to the pair $(C, \overline{f:o})$ of its class name C and the list of its fields with corresponding objects \overline{o} ; this mapping is denoted by $o \mapsto (C, \overline{f:o})$.

The form $h[o \mapsto h(o)[f \mapsto o']]$ denotes the update of the field f of the object o with the object o'.

Channel names are mapped to queues of objects: $\mathbf{k} \mapsto \overline{\mathbf{o}}$. The heap produced by $h[\mathbf{k} \mapsto \overline{\mathbf{o}}]$ maps the channel \mathbf{k} to the queue $\overline{\mathbf{o}}$. With some abuse of notation we write $\mathbf{o} :: \overline{\mathbf{o}}$ and $\overline{\mathbf{o}} :: \mathbf{o}$ to denote the queue whose first and last element is \mathbf{o} , respectively.

Heap membership for object identifiers and channels is checked using standard set notation, by identifying *h* with its domain, we can also write $o \in h$, and $k \in h$.

The values that can result from normal termination are parallel threads of fully evaluated objects.

In the reduction rules we make use of the special *channel addition* operation $\{..., \}$, and of the *continuation replacement* operation $\lfloor.../cont\rfloor$ (their formal definitions are in Figure 7, where $\{\overline{C \Rightarrow e}\}$ is short for $\{C_1 \Rightarrow e_1 \parallel C_2 \Rightarrow e_2\}$). We denote by e(k) the source expression e in which all occurrences of receive, send, and delegation expressions which *are not* within the co-body of a session request are extended, so that they explicitly mention the channel k they will use (remember that channel names are not written by the programmer). Also, we denote by $e\lfloor e'/cont\rfloor$ the expression e in which all occurrences of a session request or within the body of a send/receive loop, are replaced by e', thus preserving the correct nested structure of while expressions. For example $recC(x)\{C_1 \Rightarrow x \parallel C_2 \Rightarrow cont\}(k) \lfloor e'/cont \rfloor = k.recC(x)\{C_1 \Rightarrow x \parallel C_2 \Rightarrow e'\}$.

The reduction is a relation between pairs of threads and heaps:

$$P, h \longrightarrow P', h'$$

Reduction rules use evaluation contexts (based on runtime syntax) that capture the notion of the "next subexpression to be reduced":

$$\begin{split} \mathscr{E} &::= [-] \mid \mathscr{E}; \mathsf{e} \mid \mathscr{E}.\mathsf{f} \mid \mathscr{E}.\mathsf{f} := \mathsf{e} \mid \mathsf{o}.\mathsf{f} := \mathscr{E} \mid \mathscr{E}.\mathsf{s} \{\mathsf{e}\} \mid \\ & \mathscr{E} \bullet \mathsf{s} \{\mathsf{k}\} \mid \mathsf{k}.\mathsf{sendC}(\mathscr{E}) \{\mathsf{C}_1 \Rightarrow \mathsf{e}_1 \, \big\| \, \mathsf{C}_2 \Rightarrow \mathsf{e}_2 \} \end{split}$$

The explicit mention of the evaluation context is needed in rule SESSREQ-R (Figure 8), in which a new thread is generated in parallel with the evaluation context.

Reduction rules are in Figure 8. Rule PAR-R models the execution of parallel threads. In this rule 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 successive four rules define the execution of standard object-oriented constructions.

Rule SESSREQ-R models the connection between the co-body e of a session request $o.s \{e\}$ and the body e' of the session s, in the class of the object o. This connection is

$$\frac{\begin{array}{c} PAR-R \\ e,h \longrightarrow P,h' \\ \hline e \mid\mid P_{1},h \longrightarrow P \mid\mid P_{1},h' \end{array}}{\operatorname{SEQ-R}} \xrightarrow{\begin{array}{c} FLD-R \\ \mathscr{E}[\mathsf{o};\mathsf{e}],h \longrightarrow \mathscr{E}[\mathsf{e}],h \end{array}} \xrightarrow{\begin{array}{c} FLD-R \\ h(\mathsf{o}) = (\mathsf{C},\overline{\mathsf{f}:\mathsf{o}}) \\ \mathscr{E}[\mathsf{o}.\mathsf{f}_{i}],h \longrightarrow \mathscr{E}[\mathsf{o}_{i}],h \end{array}} \\
\frac{fields(\mathsf{C}) = \overline{\mathsf{Tf}} \quad \mathsf{o} \notin h \\ \mathscr{E}[\mathsf{o},\mathsf{f}:\mathsf{o})],h \longrightarrow \mathscr{E}[\mathsf{o}],h::[\mathsf{o} \mapsto (\mathsf{C},\overline{\mathsf{f}:\mathsf{o}})] \end{array}}{\operatorname{\mathscr{E}}[\mathsf{o}.\mathsf{f}:=\mathsf{o}'],h \longrightarrow \mathscr{E}[\mathsf{o}'],h[\mathsf{o} \mapsto h(\mathsf{o})[\mathsf{f} \mapsto \mathsf{o}']]}$$

$$\begin{split} & \frac{\text{SESSREQ-R}}{\mathscr{E}[\mathsf{o}.\mathsf{s}\,\{\mathsf{e}\}],h\longrightarrow\mathscr{E}[\mathsf{e}(\Bbbk]]\mid|\,[\mathsf{O}/\mathsf{this}]\mathsf{e}'(\check{\Bbbk}\varsigma,h[\Bbbk,\check{\Bbbk}\mapsto(\mathsf{i})])}{\mathscr{E}[\mathsf{o}.\mathsf{s}\,\{\mathsf{e}\}],h\longrightarrow\mathscr{E}[\mathsf{e}(\Bbbk]]\mid|\,|\,[\mathsf{O}/\mathsf{this}]\mathsf{e}'(\check{\Bbbk}\varsigma,h[\Bbbk,\check{\Bbbk}\mapsto(\mathsf{i})])} \xrightarrow{\mathsf{SESSDEL-R}} \underbrace{h(\mathsf{o})=(\mathsf{C},_) \qquad \mathit{sbody}(\mathsf{s},\mathsf{C})=\mathsf{e}}{\mathscr{E}[\mathsf{o}\,\mathsf{o}\,\mathsf{s}\,\{\Bbbk\}],h\longrightarrow\mathscr{E}[[\mathsf{O}/\mathsf{this}]\mathsf{e}(\check{\Bbbk}\varsigma)],h} \\ & \frac{\mathsf{SENDCASE-R}}{\mathscr{E}[\mathsf{k}.\mathsf{sendC}(\mathsf{o})\{\mathsf{C}_1\Rightarrow\mathsf{e}_1\,|\,|\,\mathsf{C}_2\Rightarrow\mathsf{e}_2\}],h\longrightarrow\mathscr{E}[\mathsf{e}_1],h[\check{\Bbbk}\mapsto\overline{\mathsf{o}}::\mathsf{o}]} \\ & \frac{\mathsf{h}(\check{\Bbbk})=\overline{\mathsf{o}} \qquad \mathsf{h}(\mathsf{o})=(\mathsf{C},_) \qquad \mathsf{C}\Downarrow\{\mathsf{C}_1,\mathsf{C}_2\}=\mathsf{C}_i}{\mathscr{E}[\mathsf{k}.\mathsf{sendC}(\mathsf{o})\{\mathsf{C}_1\Rightarrow\mathsf{e}_1\,|\,|\,\mathsf{C}_2\Rightarrow\mathsf{e}_2\}],h\longrightarrow\mathscr{E}[\mathsf{e}(\mathsf{I},\mathsf{C}_2\}=\mathsf{C}_i]} \\ & \frac{\mathsf{h}(\check{\Bbbk})=\mathsf{o}::\,\overline{\mathsf{o}} \qquad \mathsf{h}(\mathsf{o})=(\mathsf{C},_) \qquad \mathsf{C}\Downarrow\{\mathsf{C}_1,\mathsf{C}_2\}=\mathsf{C}_i}{\mathscr{E}[\mathsf{k}.\mathsf{recC}(\mathsf{x})\{\mathsf{C}_1\Rightarrow\mathsf{e}_1\,|\,|\,\mathsf{C}_2\Rightarrow\mathsf{e}_2\}],h\longrightarrow\mathscr{E}[[\mathsf{O}/\mathsf{x}]\mathsf{e}_i],h[\check{\Bbbk}\mapsto\overline{\mathsf{o}}]} \\ \\ & \mathsf{SENDWHILE-R} \\ & \mathscr{E}[\mathsf{k}.\mathsf{sendW}(\mathsf{e})\{\mathsf{C}_1\Rightarrow\mathsf{e}_1\,|\,|\,\mathsf{C}_2\Rightarrow\mathsf{e}_2\}],h\longrightarrow\mathscr{E}[\mathsf{k}.\mathsf{sendC}(\mathsf{e})\{\mathsf{C}_1\Rightarrow\mathsf{e}_1\,|\,|\,\mathsf{C}_2\Rightarrow\mathsf{e}_2\}],h \\ & \text{where} \ \mathsf{e}_i'=\lfloor\mathsf{e}_i/\mathsf{cont}\,|\,\mathsf{k}.\mathsf{sendW}(\mathsf{e})\{\mathsf{C}_1\Rightarrow\mathsf{e}_1\,|\,|\,\mathsf{C}_2\Rightarrow\mathsf{e}_2\}/\mathsf{cont} \\ \\ & \mathsf{RECEIVEWHILE-R} \\ & \mathsf{RECEIVEWHILE-R} \\ \end{aligned}$$

$$\begin{aligned} \mathscr{E}[k.\texttt{recW}(\texttt{x})\{\texttt{C}_1 \Rightarrow \texttt{e}_1 \, \|\, \texttt{C}_2 \Rightarrow \texttt{e}_2\}], h & \longrightarrow \mathscr{E}[k.\texttt{recC}(\texttt{x})\{\texttt{C}_1 \Rightarrow \texttt{e}_1' \, \|\, \texttt{C}_2 \Rightarrow \texttt{e}_2'\}], h \\ & \text{where } \texttt{e}_i' = \lfloor\texttt{e}_i/\texttt{cont} \rfloor\texttt{k}.\texttt{recW}(\texttt{x})\{\texttt{C}_1 \Rightarrow \texttt{e}_1 \, \|\, \texttt{C}_2 \Rightarrow \texttt{e}_2\}/\texttt{cont} \end{aligned}$$

Fig. 8. Reduction Rules

established through a pair of fresh channels k, \tilde{k} . For this purpose the expression o.s {e} reduces, in the same context, to its own co-body e(k) and in parallel, outside the context, it spawns the body $[0/this]e'(\tilde{k})$ of the called session. The explicit substitution of k in e and of \tilde{k} in e' ensures that the communication is on the fresh dual channels k and \tilde{k} . Thus, an object can serve *any number* of session requests. For example,

$$\begin{array}{l} \text{o.s}\{\texttt{sendC}(5)\{C_1 \Rightarrow e_1 \mid \mid C_2 \Rightarrow e_2\}\}; \texttt{new } C(\) \longrightarrow \\ \texttt{k.sendC}(5)\{C_1 \Rightarrow e_1 \wr \texttt{k} \rbrace \mid \mid C_2 \Rightarrow e_2 \wr \texttt{k} \rbrace\}; \texttt{new } C(\) \mid \mid \\ \tilde{\texttt{k.recC}}(\texttt{x})\{C_1' \Rightarrow [^{o}/\texttt{this}]e_1' \wr \tilde{\texttt{k}} \rbrace \mid \mid C_2' \Rightarrow [^{o}/\texttt{this}]e_2' \wr \tilde{\texttt{k}} \rbrace \end{array}$$

if $recC(x)\{C'_1 \Rightarrow e'_1 \mid C'_2 \Rightarrow e'_2\}$ is the body of session s in the class of the object o. Notice that there is no ambiguity in this rule, since

$$\begin{split} (k.\texttt{sendC}(5)\{C_1 \Rightarrow e_1 \wr k \rbrace \, \| \, C_2 \Rightarrow e_2 \wr k \rbrace \, \| \\ \tilde{k}.\texttt{recC}(x)\{C_1' \Rightarrow [\texttt{0}/\texttt{this}]e_1' \wr \tilde{k} \rbrace \, \| \, C_2' \Rightarrow [\texttt{0}/\texttt{this}]e_2' \wr \tilde{k} \rbrace \}); \texttt{new C}(\) \end{split}$$

is not a thread according to the syntax of $\mathscr{F}SAM^{\vee}$.

Rule SESSDEL-R replaces the session delegation $o \circ s \{k\}$ by [o/this]e(k), where e is the body of the session s, in the class of the object o. This allows a part of the communication to be delegated 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. For example $o \bullet s \{k\} \longrightarrow k.recC(x) \{C_1 \Rightarrow [o/this]e_1(k) \| C_2 \Rightarrow [o/this]e_2(k) \}$

if recC(x){ $C_1 \Rightarrow e_1 || C_2 \Rightarrow e_2$ } is the body of session s in the class of the object o.

The communication rule for sendC, SENDCASE-R, puts the object o, *i.e.* the result of evaluating the expression e, in the queue associated to the dual channel \tilde{k} of the communication channel k. The computation then proceeds with the expression e_i , if $C_1 \neq C_2$ and C_i is the smallest class in $\{C_1, C_2\}$ to which the object o belongs. Otherwise, if $C_1 = C_2$ and o belongs both to C_1 and to C_2 , then the computation proceeds with e_1^2 . This is given by the condition $h(o) = (C, _)$ and by the following definition of $C \Downarrow \{C_1, C_2\} = C_i$, using the subtyping relation (Figure 6):

$$C \Downarrow \{C_1, C_2\} = \begin{cases} C_i & \text{if } C <: C_i \text{ and } C <: C_j \text{ with } i \neq j \text{ imply } C_j \not<: C_i, \\ C_1 & \text{if } C <: C_1 = C_2, \\ \bot & \text{otherwise.} \end{cases}$$

Dually the receive communication rule takes an object o from the queue associated to channel k and returns the expression $[o/x]e_i$, if $h(o) = (C, _)$ and $C \Downarrow \{C_1, C_2\} = C_i$.

In rules SENDCASE-R and RECEIVECASE-R it is understood that the transition cannot fire if $C \Downarrow \{C_1, C_2\} = \bot$. However we will see that $C \Downarrow \{C_1, C_2\}$ is always defined in well-typed expressions.

Rules SENDWHILE-R and RECEIVEWHILE-R simply realize the repetition using the case communication expressions. Note that $sendW(\mathscr{E})\{e_1 \Rightarrow C_1 \mid e_2 \Rightarrow C_2\}$ is not an evaluation context, since we do not want to reduce the expression which controls the loop before the application of rule SENDWHILE-R, in which the sendW expression is unfolded.

Only communication and delegation expressions containing explicit channels can be reduced. So, for example, $s endC(o) \{e\}$ and $o \bullet s \{\}$ are stuck; however, as we will see in Subsection 6.1, the latter cannot be typed and the former is not an initial expression (type soundness is only guaranteed for initial expressions).

6 Typing

Session types, ranged over by t, describe the communications that take place during a session. The syntax of session types is in Figure 9, where we use \dagger as a symbol that stands for either ! or ?. By ε we denote the *empty* communication, and the *concatenation* t₁.t₂ expresses the communications in t₁ followed by those in t₂. The session type ε is the neutral element of concatenation, so that $\varepsilon.t = t = t.\varepsilon$ for all t.

The types $\{C_1 \Rightarrow t_1 \mid C_2 \Rightarrow t_2\}$ and $\{C_1 \Rightarrow t_1 \mid C_2 \Rightarrow t_2\}$ express the sending and the receiving of an object, respectively: depending on the class of this object the communication will proceed with one of the t_i . In $\mu \alpha$. $\{C_1 \Rightarrow t_1 \mid C_2 \Rightarrow t_2\}$ the *session type variable* α can occur inside t_i with the usual meaning of representing the whole session type. We consider recursive session types modulo fold/unfold: i.e., $\mu \alpha.t =$

² In this particular case, there is no other motivation for selecting the smallest index but to avoid introducing non-deterministic choices. From this point of view, alternative solutions could be just as sound: for instance, the selection of the greatest index or linguistic restrictions on the expressions e_i , e.g., the condition $e_1 = e_2$ whenever $C_1 = C_2$.

$$\begin{array}{ll} \dagger & ::= ! \mid ? \\ t & ::= \epsilon \mid \alpha \mid \odot \mid \dagger \{ C \Rightarrow t \mid C \Rightarrow t \} \mid \mu \alpha . \dagger \{ C \Rightarrow t \mid C \Rightarrow t \} \mid t.t \end{array} \\ \end{array} \qquad \begin{array}{l} \text{direction} \\ \text{session type} \end{array}$$

Fig. 9. Syntax of Session Types

 $[\mu \alpha.t/\alpha]t$. So we equate $\mu \alpha.\dagger \{C_1 \Rightarrow t_1 || C_2 \Rightarrow t_2\}$ to $\dagger \{C_1 \Rightarrow t_1 || C_2 \Rightarrow t_2\}$ when α does not occur in $\dagger \{C_1 \Rightarrow t_1 || C_2 \Rightarrow t_2\}$.

The type \odot is used only as session type for 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 SENDW-T and RECEIVEW-T in Figure 11).

We say that a session type is *closed* if it does not contain occurrences of free session type variables and of \odot . Therefore, each closed session type has one of the following shapes:

- ε;

- $\mu \alpha$. $\dagger \{C_1 \Rightarrow t_1 \parallel C_2 \Rightarrow t_2\}$ or $\dagger \{C_1 \Rightarrow t_1 \parallel C_2 \Rightarrow t_2\}$;

or a concatenation of the session types above. For simplicity we will use in definitions unfolded recursive types whenever possible.

6.1 Typing of Channel Free Expressions

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 these expressions *channel free expressions*. The term environments therefore will contain also type assignments to object identifiers. This permits a simpler formulation of the runtime typing rules, as we will see in next subsection.

The typing judgement has the shape

$$\Gamma \vdash \mathsf{e} : \mathsf{T}$$
; t

where Γ is a term environment, which maps this, cont, variables and objects to types, and t represents the session type of the (implicit) active channel.

In order to allow (possible) multiple occurrences of a variable or cont with different types inside an expression, we define the following "update" operation on $\Gamma(z \text{ ranges over this, cont, term variables, and object identifiers):$

$$\Gamma(z:T)(z') = \begin{cases} T & \text{if } z' = z \\ \Gamma(z') & \text{otherwise.} \end{cases}$$

Thus, the operation $\Gamma(z:T)$ has the effect of adding z:T to Γ , but after deleting a declaration of z from Γ (if there is one). This will avoid checking well-formedness of term environments and does not require an explicit weakening rule to add an assumption on the cont variable (when typing nested while communication expressions).

To assure a safe communication between two threads we must require their session types to be *dual*, i.e., that each send will correspond to a receive and vice versa. The duality is then the symmetric relation generated by the rules of Figure 10, in which we consider folded recursive types, otherwise the definition would not be well-founded.

$$\varepsilon \bowtie \varepsilon$$
 $\alpha \bowtie \alpha$ $\frac{\mathbf{t}_1 \bowtie \mathbf{t}'_1 \quad \mathbf{t}_2 \bowtie \mathbf{t}'_2}{\mathbf{t}_1 \cdot \mathbf{t}_2 \bowtie \mathbf{t}'_1 \cdot \mathbf{t}'_2}$

$$\frac{\mathsf{C}_1 \lor \mathsf{C}_2 <: \mathsf{C}_1' \lor \mathsf{C}_2' \quad \mathsf{C}_i \Downarrow \{\mathsf{C}_1', \mathsf{C}_2'\} = \mathsf{C}_j' \Rightarrow \mathsf{t}_i \bowtie \mathsf{t}_j' \quad \mathsf{C}_l' \Downarrow \{\mathsf{C}_1, \mathsf{C}_2\} = \mathsf{C}_k \Rightarrow \mathsf{t}_k \bowtie \mathsf{t}_l'}{\mu \alpha .! \{\mathsf{C}_1 \Rightarrow \mathsf{t}_1 \, \| \, \mathsf{C}_2 \Rightarrow \mathsf{t}_2\} \bowtie \mu \alpha .? \{\mathsf{C}_1' \Rightarrow \mathsf{t}_1' \, \| \, \mathsf{C}_2' \Rightarrow \mathsf{t}_2'\}}$$

Fig. 10. Duality Relation

The exchanged values must also be of one of the classes expected by the receiver. All possible choices on the basis of the class of the exchanged value must continue with session types which are dual of each other. For this reason we have to perform checks on the type of the exchanged values in both directions:

- for any sent value of type C_i such that $C_i \Downarrow \{C'_1, C'_2\} = C'_j$ for some $1 \le j \le 2$ we require $t_i \bowtie t'_j$;
- for any received value of type C'_l such that $C'_l \Downarrow \{C_1, C_2\} = C_k$ for some $1 \le k \le 2$ we require $t_k \bowtie t'_l$.

For instance, let us consider the session types $!{Shape \Rightarrow t_1 || String \Rightarrow t_2}$ and $?{Triangle \Rightarrow t_3 || Object \Rightarrow t_4}$ where Triangle <: Shape. At run time a Triangle can be sent as a Shape, thus the types t_1 and t_3 have to be dual. Notice that, thanks to the absence of generics we can be more flexible w.r.t. [6]: the types used in the choices (actually their union) of the send can be subtypes of the ones expected (in the dual receive).

Typing rules for channel free expressions are in Figure 11. For the sake of simplicity in rule NEWC-T we require that the initialisation of an object does not involve communications. Notice that in rule SEQ-T we use session type concatenation to represent that first the communications in e_1 and then those in e_2 are performed.

The rule for session request SESSREQ-T relies on the duality relation (Figure 10) to assure that all the bodies of the session s in the classes which build the union type T and the co-body e' of the request will communicate properly. In typing session delegation (rule SESSDEL-T) we take into account that the whole expression will be replaced by the session body defined in the class of the expression to which the session is delegated (cf. the reduction rule SESSDEL-R, Figure 8). Notice that the condition $stype(s,T)=\{t'\}$ does not imply T be one class, but only that all definitions of s in the classes which build T have the same session types. If a session has session type ε , then it is meaningless to use it in a delegation, while it is sensible to use it in a request. For this reason we require $t' \neq \varepsilon$ in rule SESSDEL-T.

Rules SENDC-T and RECEIVEC-T require all possible alternative expressions to have the same type T, but they can implement different communication sequences t_i . Rule SENDC-T prescribes that the class type of e is the union type of the classes used in the choice. Without union types the typing rule for the same construct in [18] was much more demanding and less clear. The typing rules for the while communication expressions are similar, but they also discharge the assumption on cont and replace the occurrences of \odot in session types by a fresh variable α which will be bound by μ . In rule SENDW-T typing e with session type ε prevents e from containing occurrences

Over T

XIOM-T	CONT-T	SUB-I
$\Gamma \vdash z : \Gamma(z) $; $\varepsilon \qquad z \neq cont$	$\Gamma \vdash \text{cont} : \Gamma(\text{cont})$	
NEWC T	EID	$I' \vdash e: T'$;t
$fields(C) = \overline{Tf} \qquad \Gamma \vdash e$	r_{LD} .	Γ⊢e:T;°t
$\Gamma \vdash \text{new } C(\overline{e}) : C $	$\overline{\varepsilon}$ $\Gamma \vdash \epsilon$	$e.f: ftype_r(f,T)$; t
SEQ-T	FLDASS-T	
$\Gamma \vdash e : T$; t $\Gamma \vdash e'$	$: T' $ $\mathfrak{g} t' \qquad \Gamma \vdash e : T \mathfrak{g} t$	$\Gamma \vdash e' : ftype_w(\mathtt{f}, \mathtt{T}) \mathring{s} \mathtt{t}'$
$\Gamma \vdash e;e':T'$ \mathfrak{s} t.t	$\Gamma \vdash e.f:=$	$= e' : ftype_w(f,T)$;t.t'
SESSREO-T		
Γ⊢e:T;t	$\Gamma \vdash e' : T' \ ; t' \qquad t' \bowtie t''$	$\forall t'' \in stype(s,T)$
,	$\Gamma \vdash e.s\{e'\}:T'st$	
SessDel-T		
$\Gamma \vdash e : T$ st s	$type(s,T) = \{t'\}$ $t' \neq \varepsilon$	rtype(s,T) = T'
	$\Gamma \vdash e \bullet s \{\} : T' \mathrm{\r{s}} t.t'$	
SENDC-T		
$\Gamma \vdash$	$e: C_1 \lor C_2 \mathrm{\re} t \qquad \Gamma \vdash e_i:$	T g t _i
$\Gamma \vdash sendC(e)\{C\}$	$\mathbf{h}_1 \Rightarrow \mathbf{e}_1 [] \mathbf{C}_2 \Rightarrow \mathbf{e}_2 \} : \mathbf{T} \mathring{\mathbf{g}} \mathbf{t}.! \{ \mathbf{e}_1 \} $	$C_1 \Rightarrow t_1 [] C_2 \Rightarrow t_2 \}$
RECEIVEC-T		
	$\Gamma(\mathbf{x}:\mathbf{C}_i) \vdash \mathbf{e}_i:\mathbf{T}\mathbf{\mathring{g}}\mathbf{t}_i$	
$\Gamma \vdash \texttt{recC}(\texttt{x}) \{\texttt{C}$	$\mathbf{f}_1 \Rightarrow \mathbf{e}_1 [] \mathbf{C}_2 \Rightarrow \mathbf{e}_2 \} : \mathbf{T}_9^\circ ? \{ \mathbf{C}_1 \}$	$\Rightarrow t_1 [\!] C_2 \Rightarrow t_2 \}$
SendW-T		
$\Gamma \vdash e : C_1 \lor C_2 \mathring{s} \varepsilon$	$\Gamma(\texttt{cont}:\texttt{T}) \vdash \texttt{e}_i:\texttt{T}\texttt{g}\texttt{t}_i$	α fresh in t ₁ , t ₂
$\Gamma \vdash send \mathbb{W}(e) \{ C_1 \Rightarrow e_1 \mid $	$[] C_2 \Rightarrow e_2 \} : T \mu \alpha . ! \{ C_1 \Rightarrow$	$[\alpha/ \odot] \mathtt{t}_1] \mathtt{C}_2 \Rightarrow [\alpha/ \odot] \mathtt{t}_2 \}$
RECEIVEW-T		
Γ(cont : T	$)(\mathbf{x}: C_i) \vdash \mathbf{e}_i : \mathbf{T} \circ \mathbf{t}_i \qquad \alpha \mathbf{f}$	fresh in t_1, t_2
$\Gamma \vdash \texttt{recW}(\texttt{x})\{\texttt{C}_1 \Rightarrow \texttt{e}_1 [$	$C_2 \Rightarrow e_2$: T $\circ \mu \alpha$.?{ $C_1 \Rightarrow$	$[\alpha/ \odot] \mathtt{t}_1 [\!] \mathtt{C}_2 \Rightarrow [\alpha/ \odot] \mathtt{t}_2 \}$

Fig. 11. Typing Rules for Channel Free Expressions

of communications and cont³. Notice that requiring that both branch expressions in a choice operation have the same union type T does not imply that we require them to be of the same class: in fact, T can be a proper union type. For instance, we may have that $\Gamma \vdash e_1 : T_1 \circ t_1$ and $\Gamma \vdash e_2 : T_2 \circ t_2$; by subsumption (rule SUB-T) we also have that $\Gamma \vdash e_1 : T_1 \vee T_2 \circ t_1$ and $\Gamma \vdash e_2 : T_1 \vee T_2 \circ t_2$.

Figure 12 defines well-formed class tables. Rule SESS-WF type checks the session bodies with respect to the current class C taking as term environment the association between this and C. Notice that \odot has no dual type, so sessions whose bodies would be typed with types containing \odot would be useless. This justifies the condition that t must be closed in rule SESS-WF.

A last remark is that, since no typing rule generates free session type variables, then all session types in typing judgements are closed unless they contain occurrences of \odot .

³ Note that this typing allows e to contain session requests, since the execution of these requests will use different channels to communicate.

$$\frac{SESS-WF}{\{ this: C \} \vdash e: T \ \ \ t \ \ is \ closed}}{Tts \{ e \} \ ok \ in \ C} \qquad \qquad \frac{CLASS-WF}{D \ \ ok \ \ \overline{S} \ ok \ in \ C}}{class \ C \ \ \ \ \overline{S} \ \ ok \ in \ C}$$

Fig. 12. Well-formed Class Tables

The rules, presented in this section, define how the type system checks that the declarative part of the program and the main part are well-typed, also with respect to session types used in declarations of sessions. However, when considering well-typed executable programs, we require that they are closed with respect to term variables and that all communication expressions occur inside session co-bodies, that is, that they are typed in the empty type environment with an empty session type. Namely, an *initial* expression e is such that $\emptyset \vdash e : T \ \varepsilon$ for some T. It is easy to verify that the set of initial expressions is the set of closed and well-typed user expression. For example, let us consider the stuck expressions sendC(o){e} and o • s {}:

- sendC(o){e} is well-typed but its session type is not empty,

 $- \circ \circ s$ {} is not well-typed.

We conclude this subsection by comparing $\mathscr{F}SAM^{\vee}$ with FJ \vee , an extension of FJ with union types, proposed by Igarashi and Nagira in [29]. They define union types as in the present paper: the essential difference is that they have traditional methods instead of sessions.

The method signatures are of the shape $\overline{T} \to T$, where both the parameter types \overline{T} and the return type T are union types. The method type lookup function applied to a method name m and to a union type T gives a set of method signatures, *i.e.* all the signatures which m has in the classes which build T. This is similar to our *stype* function, which returns a set of session types.

The rule of method call checks that the types of the parameters agree with all the signatures found by the method type lookup function for the union type of the object. Also our rule SESSREQ-T requires the session type of the co-body be dual to all the session types returned by the *stype* function.

It is easy to check that the encoding of methods by sessions sketched at the end of Section 3 extends without changes to methods with union types.

6.2 Typing of Runtime Expressions

During evaluation of well-typed programs, channel names are made explicit in send and receive expressions, as well as in session delegation. Thus, in order to show how well-typedness is preserved under evaluation, we need to define new typing rules for runtime expressions. Furthermore, in typing runtime expressions, we must take into account the session types of more than one channel: runtime expressions contain explicit channel names (used for communication) thus session types must be associated with channel names in an appropriate way. Then judgements have the form

$$\Gamma \vdash_{\mathbf{r}} \mathsf{e} : \mathsf{T} \, \mathrm{s} \, \Sigma$$

where Σ denotes a *session environment* which maps channels to session types.

		FLD-RT		SUB-RT	
AXIOM-RT $\Gamma \vdash \pi \cdot \Gamma(\pi) \circ 0$	$\pi \neq cont$	$\Gamma \vdash_{\mathbf{r}}$	e:T $ m \resize{\Sigma}$	$\Gamma \vdash_{\mathbf{r}} e : T \Sigma$	T <: T'
$I \vdash_{\mathbf{r}} \mathbf{Z} : I(\mathbf{Z}) \circ \mathbf{\emptyset} \qquad \mathbf{Z} \neq \text{cont}$	z≠cont	$\Gamma \vdash_{\mathbf{r}} \mathbf{e.f} : f$	$type_r(f,T)$; Σ	$\Gamma \vdash_{\mathbf{r}} e : T'$	ŝΣ
NEWC-RT			CONT PT		
fields(C) = Tf	$\Gamma \vdash_{\mathbf{r}} e_i : I$	Γ _i	$\Gamma(cont \cdot T)$	= cont · T ∘ ∫k · ⊜ ໄ	
$\Gamma \vdash_{\mathbf{r}} nev$	v C(ē): C ᡲØ		1 (conc. 1)	r conc. 19 [k.⊚]	
SEQ-RT		, FL	dAss-RT		
$\Gamma \vdash_{\mathtt{r}} \mathtt{e} : \mathtt{T} \ $	$\Gamma \vdash_{\tt r} {\sf e}' : {\tt T}'$	$\frac{2}{5}\Sigma'$ Γ	_re:TŝΣ	$\Gamma \vdash_{\mathbf{r}} \mathbf{e}' : ftype_w(\mathbf{f}, \mathbf{T})$	m
Γ⊢ _r e;e	$e': T' ightharpoonup \Sigma.\Sigma'$		$\Gamma \vdash_{\mathbf{r}} e.f := e'$: $ftype_w(f,T)$; Σ . Σ'	
SESS	REQ-RT				
$\Gamma \vdash_{\mathbf{r}} e$	$E:T$ \mathcal{S} Σ Γ	'⊢ e′: T′ ; t′	t′⊠t″∀t	$z'' \in stype(s, T)$	
		Γ⊢ _r e.s{e	$^{\prime}\}:\mathtt{T}^{\prime}$ $^{\circ}\varSigma$		
SESSD	EL-RT				
<i>I</i> ⊢ _r e:	$1 \stackrel{\circ}{,} \Sigma styp$	$pe(s,T) = \{t$	} t≠ε	$rtype(s,T) = T^*$	
	Г	$\vdash_{\mathbf{r}} \mathbf{e} \bullet \mathbf{s} \{\mathbf{k}\}:$	$\texttt{T}' \texttt{\$} \boldsymbol{\Sigma}.\{\texttt{k}:\texttt{t}\}$		
SENDC-R					
	I Fre: C	1 V C2 § Z	$I \vdash_{\mathbf{r}} \mathbf{e}_i \colon \mathbf{I} \not\in \{\mathbf{K}\}$.: L _i }	
∏ ⊢ _r k.ser	$\mathrm{udC}(e)\{C_1\Rightarrowe\}$	$\mathbf{e}_1 \parallel \mathbf{C}_2 \Rightarrow \mathbf{e}_2$	$: T ; \Sigma : \{k : ! \{C\}$	$_1 \Rightarrow t_1 \parallel C_2 \Rightarrow t_2 \} \}$	
RECEIVI	EC-RT		- (1)		
	1	$(\mathbf{x}:C_i)\vdash_{r}e_i$	$: T $ {k $: t_i$ }		
$\Gamma \vdash_{\mathbf{r}} \mathbf{k}.\mathbf{r}$	$ecC(x) \{C_1 \Rightarrow$	$e_1 [\!] C_2 \Rightarrow e_2$	$\}: T_{9}^{\circ} \{k:?\{C_{1}\}$	$\Rightarrow t_1 [C_2 \Rightarrow t_2 \} \}$	
SENDW-RT					
$\Gamma \vdash_{r} e : C_{1}$	$\vee C_2 \mathrm{\ress} \emptyset \qquad \Gamma$	$(\texttt{cont}:\texttt{T}) \vdash_{\texttt{r}}$	$e_i : T \circ \{k : t_i\}$	α fresh in t ₁ , t ₂	2
$\Gamma \vdash_{\mathbf{r}} \mathbf{k}.sendW(\mathbf{e})$	$C_1 \Rightarrow e_1 \ C_2 =$	$\Rightarrow e_2 \} : T $	$\mathbf{k}: \mu \alpha.! \{ C_1 \Rightarrow [$	$\alpha / \odot] t_1 \ C_2 \Rightarrow [\alpha / \alpha]$	$\mathbb{O}[t_2]$
RECEIVEW					- 2
$\Gamma($	cont : T)(x : C	$\mathbf{z}_i) \vdash_{\mathbf{r}} \mathbf{e}_i : \mathbf{T}$ \$	$\{k:t_i\}$ α for	tesh in t_1, t_2	
$\Gamma \vdash_{\mathbf{r}} \mathbf{k}.\mathbf{recW}(\mathbf{x}) \{ 0 \}$	$C_1 \Rightarrow e_1 C_2 \Rightarrow$	> e 2 } : T % {k	$: \mu\alpha.?{C_1 \Rightarrow [0]}$	$\alpha/\odot]t_1 []C_2 \Rightarrow [\alpha/\varpi]$	\mathbb{I}_{2}
· · · · · · · · · · · · · · · · · · ·	1 / -1 0 -2 /	2)·-/(· · · · · · · · · · · · · · · · · · ·		-1-711

Fig. 13. Typing Rules for Runtime Expressions

A session environment maps only a finite set of channels to session types different from ε , and all the remaining to ε . We can then represent one session environment with an infinite number of finite sets which give all the meaningful associations and some of the others. For example $\{k : t\}$ and $\{k : t, k' : \varepsilon\}$ represent the same environment. This choice avoids an explicit weakening rule for session environments. Figure 13 gives the typing rules for runtime expressions, which differ from those for channel free expressions for having session environments instead of a unique session type. For this reason we extend the *concatenation* of session types to *session environments* as follows:

$$\Sigma . \Sigma'(\mathbf{k}) = \Sigma(\mathbf{k}) . \Sigma'(\mathbf{k})$$

Notice that in rule SESSREQ-RT we are making use of the judgement $\Gamma \vdash e': T$; t', where the expression e' does not contain channels, but it can contain object identifiers. This justifies our choice of considering channel free expressions instead of user expressions in the typing rules of previous subsection. Notice also that the session environments of the branches in the communication expressions only contain the current channel as subject, since these expressions will never be reduced before the selection

has been done. In rule SENDW-RT we assume $\Gamma \vdash_{r} e : C_1 \lor C_2 \ 0$, since the evaluation of e cannot start before the sendW expression has been unfolded to a sendC.

The typing rules for runtime expressions differ from the ones for user expressions only in assigning the session type to explicit channels, not in the union type.

6.3 Type Soundness

Our type system enjoys subject reduction and assures progress $(\longrightarrow^*$ is the reflexive and transitive closure or \longrightarrow):

- If $\emptyset \vdash e_0 : T_0 \$ ε and $e_0, [] \longrightarrow P, h$, where $P \equiv e_1 || \dots || e_n$, then:
 - for each e_i we get $\Gamma \vdash_{\mathbf{r}} e_i : T_i \circ \Sigma_i$ for some Γ, T_i, Σ_i $(1 \le i \le n)$, and there is j $(1 \le j \le n)$ such that $T_j = T_0$, and;
 - either $P, h \longrightarrow P', h'$ for some P', h', or for all $i (1 \le i \le n) e_i$ is an object identifier. The runtime errors which our type system prevents are:
 - 1. the selection of a field and the request of a session which do not belong to the class of the current object;
 - 2. the creation of a pair of dual channels whose communication sequences do not perfectly match.

Proofs, more examples and discussions can be found in the extended version of this paper, available at http://www.di.unito.it/~dezani/papers/bcdgvfull.pdf.

7 Conclusion

In the present paper we showed, through the language SAM^{\vee} , how the addition of union types to an object oriented language with session types enhances flexibility.

The amalgamation of communication centred and object oriented programming, as it has been developed in [18,6] and in the present paper, can be extended in various directions. In particular we plan to integrate this approach with multi-party session communication [9] and with safe failure recovery [4].

Moreover, we want to study the extension of union and intersection to session types, following the intuition given by union/intersection of contracts in [12].

Lastly it would be interesting to integrate session primitives with name constraints as introduced in [5] in order to allow specification of Quality of Service requirements.

Acknowledgements. We thank the referees for their helpful comments. The final version of the paper improved due to their suggestions.

References

- Barbanera, F., Dezani-Ciancaglini, M., de'Liguoro, U.: Intersection and Union Types: Syntax and Semantics. Information and Computation 119, 202–230 (1995)
- Bonelli, E., Compagnoni, A.: Multipoint Session Types for a Distributed Calculus. In: Barthe, G., Fournet, C. (eds.) TGC 2007. LNCS, vol. 4912, pp. 240–256. Springer, Heidelberg (to appear, 2008)
- Bonelli, E., Compagnoni, A., Gunter, E.: Correspondence Assertions for Process Synchronization in Concurrent Communications. Journal of Functional Programming 15(2), 219–248 (2005)

- Boreale, M., Bruni, R., Caires, L., Nicola, R.D., Lanese, I., Loreti, M., Martins, F., Montanari, U., Ravara, A., Sangiorgi, D., Vasconcelos, V., Zavattaro, G.: SCC: a Service Centered Calculus. In: Bravetti, M., Núñez, M., Zavattaro, G. (eds.) WS-FM 2006. LNCS, vol. 4184, pp. 38–57. Springer, Heidelberg (2006)
- Buscemi, M., Montanari, U.: CC-Pi: A Constraint-Based Language for Specifying Service Level Agreements. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 18–32. Springer, Heidelberg (2007)
- Capecchi, S., Coppo, M., Dezani-Ciancaglini, M., Drossopoulou, S., Giachino, E.: Amalgamating Sessions and Methods in Object Oriented Languages with Generics (submitted, 2007)
- Carbone, M., Honda, K., Yoshida, N.: A Calculus of Global Interaction Based on Session Types. In: Fernández, M., Kirchner, C. (eds.) SecReT 2006. ENTCS, vol. 171(3), pp. 127– 151. Elsevier, Amsterdam (2007)
- Carbone, M., Honda, K., Yoshida, N.: Structured Communication-Centred Programming for Web Services. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 2–17. Springer, Heidelberg (2007)
- Carbone, M., Honda, K., Yoshida, N.: Multiparty Asynchronous Session Types. In: Necula, G.C., Wadler, P. (eds.) POPL 2008, pp. 273–284. ACM Press, New York (2008)
- 10. Castagna, G., De Nicola, R., Varacca, D.: Semantic Subtyping for the π -calculus. In: Panangaden, P. (ed.) LICS 2005, pp. 92–101. IEEE Computer Society Press, Los Alamitos (2005)
- Castagna, G., Frisch, A.: A Gentle Introduction to Semantic Subtyping. In: Barahona, P., Felty, A.P. (eds.) PPDP 2005, pp. 198–199. ACM Press, New York (2005)
- Castagna, G., Gesbert, N., Padovani, L.: A Theory of Contracts for Web Services. In: Necula, G.C., Wadler, P. (eds.) POPL 2008, pp. 261–272. ACM Press, New York (2008)
- Coppo, M., Dezani-Ciancaglini, M., Yoshida, N.: Asynchronous Session Types and Progress for Object-Oriented Languages. In: Bonsangue, M.M., Johnsen, E.B. (eds.) FMOODS 2007. LNCS, vol. 4468, pp. 1–31. Springer, Heidelberg (2007)
- Dezani-Ciancaglini, M., de' Liguoro, U., Yoshida, N.: On Progress for Structured Communications. In: Barthe, G., Fournet, C. (eds.) TGC 2007. LNCS, vol. 4912, pp. 257–275. Springer, Heidelberg (2008)
- Dezani-Ciancaglini, M., Drossopoulou, S., Giachino, E., Yoshida, N.: Bounded Session Types for Object-Oriented Languages. In: de Boer, F., Bonsangue, M. (eds.) FMCO 2006. LNCS, vol. 4709, pp. 207–245. Springer, Heidelberg (2007)
- Dezani-Ciancaglini, M., Mostrous, D., Yoshida, N., Drossopoulou, S.: Session Types for Object-Oriented Languages. In: Thomas, D. (ed.) ECOOP 2006. LNCS, vol. 4067, pp. 328– 352. Springer, Heidelberg (2006)
- Dezani-Ciancaglini, M., Yoshida, N., Ahern, A., Drossopoulou, S.: A Distributed Object Oriented Language with Session Types. In: De Nicola, R., Sangiorgi, D. (eds.) TGC 2005. LNCS, vol. 3705, pp. 299–318. Springer, Heidelberg (2005)
- Drossopoulou, S., Dezani-Ciancaglini, M., Coppo, M.: Amalgamating the Session Types and the Object Oriented Programming Paradigms. In: MPOOL 2007 (2007), http://homepages.fh-regensburg.de/mpool/mpool07/programme.html
- Fähndrich, M., Aiken, M., Hawblitzel, C., Hodson, O., Hunt, G.C., Larus, J.R., Levi, S.: Language Support for Fast and Reliable Message-based Communication in Singularity OS. In: Zwaenepoel, W. (ed.) EuroSys 2006. ACM SIGOPS, pp. 177–190. ACM Press, New York (2006)
- Fournet, C., Laneve, C., Maranget, L., Rémy, D.: Inheritance in the Join Calculus. In: Kapoor, S., Prasad, S. (eds.) FST TCS 2000. LNCS, vol. 1974, pp. 397–408. Springer, Heidelberg (2000)
- Gapeyev, V., Pierce, B.C.: Regular Object Types. In: Cardelli, L. (ed.) ECOOP 2003. LNCS, vol. 2743, pp. 151–175. Springer, Heidelberg (2003)

- Garralda, P., Compagnoni, A., Dezani-Ciancaglini, M.: BASS: Boxed Ambients with Safe Sessions. In: Maher, M. (ed.) PPDP 2006, pp. 61–72. ACM Press, New York (2006)
- 23. Gay, S.: Bounded Polymorphism in Session Types. In: MSCS (to appear, 2008)
- 24. Gay, S., Hole, M.: Subtyping for Session Types in the Pi-Calculus. Acta Informatica 42(2/3), 191–225 (2005)
- Gay, S., Vasconcelos, V.T., Ravara, A.: Session Types for Inter-Process Communication. TR 2003–133, Department of Computing, University of Glasgow (2003)
- Honda, K.: Types for Dyadic Interaction. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 509–523. Springer, Heidelberg (1993)
- Honda, K., Vasconcelos, V.T., Kubo, M.: Language Primitives and Type Disciplines for Structured Communication-based Programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 22–138. Springer, Heidelberg (1998)
- Honda, K., Yoshida, N., Carbone, M.: Web Services, Mobile Processes and Types. EATCS Bulletin 91, 160–188 (2007)
- Igarashi, A., Nagira, H.: Union Types for Object Oriented Programming. Journal of Object Technology 6(2), 31–52 (2007)
- Igarashi, A., Pierce, B.C., Wadler, P.: Featherweight Java: a Minimal Core Calculus for Java and GJ. ACM TOPLAS 23(3), 396–450 (2001)
- 31. Sparkes, S.: Conversation with Steve Ross-Talbot. ACM Queue 4(2), 14-23 (2006)
- Takeuchi, K., Honda, K., Kubo, M.: An Interaction-based Language and its Typing System. In: Halatsis, C., Philokyprou, G., Maritsas, D., Theodoridis, S. (eds.) PARLE 1994. LNCS, vol. 817, pp. 398–413. Springer, Heidelberg (1994)
- Vallecillo, A., Vasconcelos, V.T., Ravara, A.: Typing the Behavior of Objects and Components using Session Types. In: Brogi, A., Jacquet, J.-M. (eds.) FOCLASA 2002. ENTCS, vol. 68(3), pp. 439–456. Elsevier, Amsterdam (2002)
- Vasconcelos, V.T., Gay, S., Ravara, A.: Typechecking a Multithreaded Functional Language with Session Types. Theorical Computer Science 368(1-2), 64–87 (2006)
- 35. Web Services Choreography Working Group. Web Services Choreography Description Language (2002), http://www.w3.org/2002/ws/chor/