

# Dynamic Interfaces

Simon J. Gay

Department of Computing Science  
University of Glasgow, UK  
simon@dcs.gla.ac.uk

António Ravara

SQIG at Instituto de Telecomunicações  
Department of Mathematics, IST,  
Technical University of Lisbon, Portugal  
amar@math.ist.utl.pt

Vasco T. Vasconcelos

Departamento de Informática  
Faculdade de Ciências da  
Universidade de Lisboa, Portugal  
vv@di.fc.ul.pt

## Abstract

We define a small class-based object-oriented language in which the availability of methods depends on an object's abstract state: objects' interfaces are *dynamic*. Each class has a *session type* which provides a global specification of the availability of methods in each state. A key feature is that the abstract state of an object may depend on the result of a method whose return type is an enumeration. Static typing guarantees that methods are only called when they are available. We present both a type *system*, in which the typing of a method specifies pre- and post-conditions for its object's state, and a typechecking *algorithm*, which infers the pre- and post-conditions from the session type, and prove type safety results. Inheritance is included; a subtyping relation on session types, related to that found in previous literature, characterizes the relationship between method availability in a subclass and in its superclass. We illustrate the language and its type system with example based on a Java-style iterator and a hierarchy of classes for accessing files, and conclude by outlining several ways in which our theory can be extended towards more practical languages.

**Categories and Subject Descriptors** D.3.3 [Language Constructs and Features]: Classes and objects; D.3.2 [Language Classifications]: Object-oriented languages; D.3.1 [Formal Definitions and Theory]; F.3.2 [Semantics of Programming Languages]: Operational semantics; F.3.3 [Studies of Program Constructs]: Type structure; D.1.5 [Object-oriented Programming]

**General Terms** Languages, Theory, Verification

**Keywords** Session types, object-oriented calculus, non-uniform method availability

## 1. Introduction

Standard class-based object-oriented languages present the programmer with the following view: an object is declared to belong to a certain type, and the type defines various methods; it is therefore possible at any time to call any of the methods described in the type. However, there are often semantic reasons why it is not appropriate to call a particular method when the object is in a particular internal state. For example: with a stack, one should not attempt to pop if the stack is empty; with a finite buffer, one should not attempt

to write if the buffer is full; with a file, one should not attempt to read data until the file has been opened. We will refer to the set of available methods as the *interface* of an object, and use the term *dynamic interfaces* in connection with objects for which method availability depends on state. Objects with dynamic interfaces are also referred to in the literature as *non-uniform objects*, although usually in a concurrent setting.

Consider the `java.util.Iterator` interface.

```
interface Iterator {
  boolean hasNext()
  Object next()
  void remove()
}
```

One correct pattern for using an iterator it is as follows:

```
while (it.hasNext())
  { ... it.next() ... }
```

while common programming errors, familiar to any teacher, include failing to call `hasNext()`:

```
for (int i = 0; i < 5; i++) {
  ... it.next() ... }
```

calling `next()` too often:

```
while (it.hasNext())
  if (it.next().equals(x))
    y.add(it.next());
```

and omitting an initial call of `hasNext()`:

```
b.append(it.next());
while (it.hasNext())
  b.append(", ").append(it.next());
```

There are two problems with the above interface: a) it provides no clue on how the interface should be used; in fact it is not obvious from the English text that accompanies `java.util.Iterator` when one is supposed to call method `remove`; b) all of the code snippets above compile, errors being caught only at runtime, if ever.

We propose to discipline the order of method calls by annotating the above interface with a *session type*, inspired by work on type-theoretic specifications of communication protocols (40; 25), as follows.

```
session Init
where Init = &{hasNext: Result}
      Result =  $\oplus$ {true: Next,
                false: end}
      Next = &{next: &{hasNext: Result,
                    remove: Init},
            hasNext: Result}
```

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FOOL 2009 24 January, 2009, Savannah, Georgia, USA.  
Copyright © 2009 ACM [to be supplied]...\$5.00

The session type defines certain abstract states of iterator objects and specifies which methods are available in each state. For example, in state `Init` only the method `hasNext()` is available, and after calling it the abstract state becomes `Result`. The form of this state indicates that the method `hasNext()` returns a result of type `boolean` and that the subsequent state depends on the result. In general, `&` indicates available methods and  $\oplus$  indicates possible results. In state `Next`, the method `next()` is available, and after calling it, the method `remove()` is also available. The keyword `end` is an abbreviation for an empty `&`, indicating that no methods are available, hence that the object cannot be used further and may be subject to garbage collection. The session type captures the specification of an iterator: `hasNext()` must be called in order to find out whether or not `next()` can be called, and `remove()` can be called at most once for each call of `next()`. Our type system makes sure that, not only methods are called by the specified order, but also that client code actually tests methods' results and proceeds accordingly, rendering the above programming errors untypable.

In the present paper we propose, in the sequential setting, a type system to support static checking of correctness of method calls in the presence of objects with dynamic interfaces. The key features of our approach are as follows.

- Each class declares a *session type* which provides an abstract view of the allowed sequences of method calls. In the simplest case, a session type defines a directed graph of abstract state transitions labelled by method names.
- A session type can also specify that the abstract state after a method call depends on the result of the call, where this result comes from an enumerated type. In this case the caller must perform a case-analysis on the result before calling further methods.
- We allow inheritance between classes, characterizing the conditions for inheritance by means of a subtyping relation on session types.
- We formalize the operational semantics and typing rules of a core language with these features, enabling us to prove a type safety result. Objects with dynamic interfaces are handled linearly in order to avoid aliasing problems.
- We define a typechecking algorithm, whose success guarantees not only type safety but also consistency between method definitions and session types, so that every sequence of methods calls in the session type is realizable in a typable program.

There is a substantial literature of related work, which we discuss in more detail in Section 6. Type systems for non-uniform concurrent objects have been proposed by several authors including Nierstrasz (31), Ravara *et al.* (38; 39) and Puntigam and Peter (35; 36). The type systems of the imperative languages *Cyclone* (21; 22), *Vault* (9; 16) and *Fugue* (17; 10) address similar issues in sequential programming. The related topic of type-theoretic specifications of protocols on communication channels has been studied for concurrent object-oriented languages by Dezani-Ciancaglini *et al.* (11; 12; 13; 14) and implemented in the *Sing#* language (15).

## Contributions

Our work makes the following new contributions. In contrast with *Cyclone* and *Vault*, we define an object-oriented language with inheritance. In contrast with *Sing#* and other work on session types for object-oriented languages, we consider objects with dynamic interfaces in a setting more general than communication channels. In contrast with *Fugue*, we use a session type as a global specification of method availability, instead of pre- and post-conditions on methods. In contrast with work on non-uniform concurrent objects, we work with a Java-like language, not a mobile process calculus;

```

enum Res {OK, NOT_FOUND, DENIED;}           1
                                           2
enum Bool {FALSE, TRUE;}                   3
                                           4
interface FileReadToEnd {                   5
  session Init                               6
  where Init = &{open:  $\oplus$ {Res.OK: Open,    7
                          Res.NOT_FOUND: end, 8
                          Res.DENIED: end}} 9
      Open = &{eof:  $\oplus$ {Bool.TRUE: Close,   10
                       Bool.FALSE: Read}} 11
      Read = &{read: Open}                   12
      Close = &{close: end}                  13
                                           14
  Res open()                                 15
                                           16
  Bool eof()                                 17
                                           18
  String read()                              19
                                           20
  Null close()                               21
}                                             22

```

**Figure 1.** The interface of a file that must be read to the end-of-file.

and we force the caller to perform a case-analysis when necessary, meaning that our session types are richer than simply sequences of available methods.

The remainder of the paper is structured as follows. In Section 2 we illustrate our system by means of a more extensive example, extending it in Section 3 to include inheritance. In Section 4 we formalize a core language and prove a type safety result. The core language requires, in addition to the session type, explicit pre- and post-conditions for each method. In Section 5 we present a type-checking algorithm which infers the pre- and post-conditions from the session types. Section 6 contains a more extensive discussion of related work, Section 7 outlines future work and Section 8 concludes.

## 2. Programming with Dynamic Interfaces

This section and the next section contain more extensive examples of programming with dynamic interfaces. From now on, we do not use the term *interface* in the technical Java sense; for us, an interface is simply a class definition without the code of the methods. Of course it includes a session type, making it dynamic.

Our next example involves the class `FileReadToEnd`, representing part of an API for using a file system. A file has a dynamic interface: it must first be opened, then can be read repeatedly, and must finally be closed. Before reading, a test for end of file must be carried out, in a way similar to the iterator. A key feature of `FileReadToEnd` is that the file cannot be closed until all of the data has been read. We will relax this restriction later. The example also contains a class `FileReader`, representing application code which uses `FileReadToEnd` to access a file and read its data into a string.

Figures 1 and 2 contain the code for the example. Figure 1 consists of three declarations. Lines 1 and 3 define enumerated types `Res` and `Bool`. Lines 5–22 define the interface `FileReadToEnd`. For technical reasons we use an enumerated type `Bool` rather than the primitive type `boolean`, and type `Null` rather than `void`. The session type is represented diagrammatically in Figure 7(a).

Our language does not include constructor methods as a special category, but the method `open` must be called first and can therefore be regarded as a constructor.

Figure 2 defines the class `FileReader`, which uses an object of class `FileReadToEnd`. The class diagram for the example is in Figure 8. This class has a session type of its own, defined on lines 2



```

class FileBoundedReader extends FileReader {
  session &{init: &{read: Final,
                    readNext: Final}}
  where ... // see Fig. 3
  @Override
  Null init() {
    f = new FileRead(); s = ""; }

  Null readFirst() {
    g = f.open();
    switch(g) {
      case Res.NOT_FOUND:
      case Res.DENIED: break;
      case Res.OK:
        b = f.eof();
        switch(b) {
          case Bool.TRUE: break;
          case Bool.FALSE:
            f.read();
            f.close();
            break;
        }
      break; } }
}

```

Figure 5. A client that reads strings from a FileReadToEnd.

```

enum KindRes restricts Res {OK;}
interface KindFileRead extends FileRead {
  session Init
  where Init = &{open: ⊕{KindRes.OK: Open}
    ... // see Fig. 5
  @Override
  KindRes open()
}

```

Figure 6. A file that can always be opened for reading.

**extends** form for classes, but instead of specifying additional members it specifies the remaining values.

The original `FileReader` can safely use a `KindFileRead`; the `NOT_FOUND` and `DENIED` branches of the `switch` statement will never be called. In the new `KindFileBoundedReader` class (not shown), the `switch` statement has just one branch, for `OK`. Had we defined `KindFileRead` and `KindBoundedReader` in isolation there would be no need to use an enumeration at all, but it is necessary for compatibility with the rest of the class hierarchy.

#### 4. A Core Language for Dynamic Interfaces

We now present a formal syntax, operational semantics and type system for a core language containing the essential features of the examples in Sections 1–3, and prove a type safety result. The main simplification is that methods must return enumerated types, which means that all session types alternate between  $\&$  and  $\oplus$  constructors. All objects have dynamic interfaces, meaning that all objects are handled linearly by the type system, whereas a practical language would also contain standard (non-dynamic) objects. We also omit method parameters, which complicate the system but don't require new conceptual techniques. The formal language only includes classes; the interfaces in Sections 2 and 3 should be interpreted as classes without method bodies.

In order to simplify the presentation of the type system and the proof of type safety, the formal language requires every method definition to be annotated with pre- and post-conditions, expressed

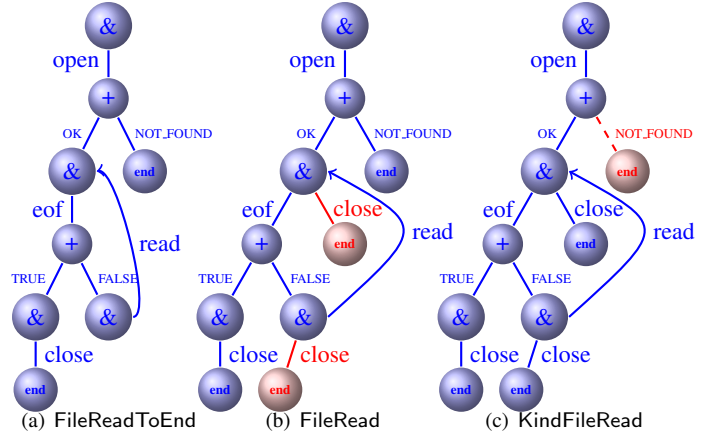


Figure 7. The diagrammatic representation of the session types for the classes in Figures 1, 5 and 6 (branch `DENIED` omitted).




Figure 8. Class diagram for the classes in the example.

Class dec	$D ::= \text{class } C \{S; \vec{f}; \vec{M}\} \mid \text{enum } E L \mid \text{class } C \text{ extends } C \{S; \vec{f}; \vec{M}\} \mid \text{enum } E \text{ restricts } E L$
Constant sets	$L ::= \{l_i\}_{i \in I}$
Method dec	$M ::= T m() \{e\}$
Values	$v ::= \text{null} \mid E.l \mid o$
Expressions	$e ::= v \mid \text{new } C() \mid o.f \mid o.f = e \mid o.f.m() \mid e; e \mid \text{while } (o.f) \{e\} \mid \text{switch } (o.f) \{\text{case } l_i : e_i\}_{i \in I}$
Types	$T ::= \text{Null} \mid E$
Session types	$S ::= \&\{m_i : S_i\}_{i \in I} \mid \oplus\{E.l_i : S_i\}_{i \in I} \mid \mu X.S \mid X$

The only object reference  $o$  available to the programmer is this.

Figure 9. Programmer's syntax

as a requirement (`req`) and a guarantee (`ens`, for “ensures”) on the type of the object on which it is called. These annotations are in the style of the Fugue system (10) but stated in terms of session types. The typechecking algorithm presented in Section 5 infers the pre- and post-conditions.

#### 4.1 Syntax

We separate the syntax into the programmer's language (Figure 9) and the extensions required by the type system and operational semantics (Figure 10). Class, enum and method declarations, including the forms for inheritance, have been illustrated by the examples.

Method dec	$M ::= \text{req } T \text{ ens } T \text{ for } T \text{ m}() \{e\}$
Values	$v ::= \dots \mid v \text{ then } f = o$
Expressions	$e ::= \dots \mid e \text{ then } f = o$
Field types	$F ::= \vec{T} \vec{f} \mid \langle E.l_i : F_i \rangle_{i \in I}$
Types	$T ::= \dots \mid C[S] \mid C[S; F] \mid T \text{ then } T \text{ f}$
Heaps	$h ::= \varepsilon \mid h :: o = C[S; V]$
Field values	$V ::= \vec{f} = \vec{v}$
States	$s ::= h; e$
Contexts	$E ::= [] \mid o.f = E \mid E; e \mid E \text{ then } f = o$

**Figure 10.** Runtime (and type system) syntax

We write  $\text{session}(C)$ ,  $\text{fields}(C)$ ,  $\text{methods}(C)$  to access the components of a class, and  $\text{constants}(E)$  for the set of values in an enum. A class declaration does not declare types for fields because they can vary at run-time. When an object is created, its fields are initialised to null. We assume that class and enum identifiers in a sequence of declarations  $\vec{D}$  are all distinct, and that method names, field names and labels in  $\vec{M}$ ,  $\vec{f}$  and  $\{l_i\}_{i \in I}$  are distinct.

In the syntax of expressions, field access, assignment and method call are only available on an object reference, not an arbitrary expression; moreover, the only object reference available to the programmer is this. All fields are private:  $\text{this}.f.g$  and  $\text{this}.f.g.m()$  are not syntactically correct. The examples in Section 2 omit this as the prefix to all field accesses, but they could be easily be inserted by the compiler. A more unusual point is that we only switch on a field, not an arbitrary expression; similarly the condition of a while loop tests a field. Section 2 illustrates the corresponding programming style, which we will discuss further during the rest of this section.

In the internal syntax (Figure 10), the expression  $e \text{ then } f = o$  and the corresponding type  $T \text{ then } U \text{ f}$  are essential to allow the enum result of a method to be stored and tested later.

Field values, heaps and states are used to define the operational semantics. An entry in the heap maps an object reference to an object:  $o = C[S; \vec{f} = \vec{v}]$ , where  $C$  is the class,  $S$  is the session type and the fields  $\vec{f}$  have values  $\vec{v}$ . The session type is used as an abstraction of the state, to indicate which methods are available. It is only needed in order to state type safety, namely that only available methods may be called. An implementation would not need that information: once typechecked, a program can be executed using a simplified version of heaps, without session types.

Session types have been discussed in relation to the example. Session type end abbreviates  $\&\{\}$ . In  $\oplus\{C.l_i : S_i\}_{i \in I}$ ,  $C$  is an enum with values  $\{l_i\}_{i \in I}$ . In the session type of a class declaration the top-level constructor, apart from recursion, must be  $\&$ .

The core language does not include named session types or the **session** and **where** clauses from the examples; we just work with recursive session type expressions of the form  $\mu X.S$ , which are required to be *contractive*, i.e. containing no subexpression of the form  $\mu X_1 \dots \mu X_n.X_1$ . We adopt the equi-recursive approach (34, Chapter 21) and regard  $\mu X.S$  and  $S\{\mu X.S/X\}$  as equivalent, using them interchangeably in any mathematical context.

A typing of the fields of an object is either a list  $\vec{T} \vec{f}$  of typed fields or a variant type over typed fields. In the variant type  $\langle C.l_i : \vec{T}_i \vec{f} \rangle_{i \in I}$ , the labels  $l_i$  come from an enum  $C$ . The type of an object,  $C[S; F]$ , consists of a class, a session type, and a typing of the fields. The type Null has the single value null.

Finally, the type system uses type environments  $\Gamma$ , which are functions assigning types to object references  $o$ ; the operational semantics uses reduction contexts  $E[\ ]$  in the style of Wright and Felleisen (41).

## 4.2 Operational Semantics

Figure 11 defines an operational semantics on configurations  $h; e$  consisting of a heap and an expression. The rules are implicitly parameterized by  $\vec{D}$ , the list of declarations constituting the program. We regard a heap  $h$  as a function from object references  $o$  to objects  $C[S; \vec{f} = \vec{v}]$ . If  $h(o) = C[S; \vec{f} = \vec{v}]$  then  $h(o).\text{class}$  means  $C$ ,  $h(o).\text{session}$  means  $S$ , and  $h(o).f_i$  means  $v_i$ . If  $h(o)$  is defined (this is an implicit hypothesis) then the notation  $h + o.f = v$  means the heap obtained by changing the value of field  $f$  in object  $o$  to  $v$ . Similarly  $h + o.\text{session} = S$ .

The rules central to our proposal—method call and switch—and their interplay with the novel then expression are better described via an example. The reduction of the expression

$$o.g = o.f.m_j(); \text{switch } (o.g) \{ \text{case } l_k : e_k \}$$

is illustrated in Figure 12. Calling  $o.f.m_j()$  nullifies  $o.f$  and introduces the delayed assignment  $\text{then } f = o'$ . Once the then expression reduces to a value, the method returns, and the then value is stored in field  $g$ . The subsequent switch deconstructs the then-value, selects the appropriate case and exercises the delayed assignment, storing  $o'$  back in field  $f$ .

R-NEW creates a new object in the heap, with an initial session type taken from the class declaration and with its fields set to null. R-FIELD extracts the value of a field from an object in the heap. Linear control of objects requires that the field be nullified. R-ASSIGN updates the value of a field. The value of the assignment, as an expression, is null; linearity means that it cannot be  $v$  as in Java.

R-WHILEF/R-WHILET define the behaviour of a while expression to be equivalent to a switch with two branches, one which terminates immediately and one which executes  $e$  and then loops. R-SEQ discards the result of the first part of a sequential composition. R-CONTEXT is the usual rule for reduction in contexts.

To complete the definition of the semantics we need to define the initial state. For a given class  $C$  and main method  $m$ , one would expect an initial state of the form  $\emptyset; \text{new } C().m()$ . Given that we can call methods on fields only, we set up a dummy class  $C'$ , with a completed session type (end), a single field  $f$ , and no methods. The initial state is then

$$o = C'[\text{end}; f = \text{null}]; (o.f = \text{new } C(); o.f.m())$$

## 4.3 Subtyping

The foundation for the theory of inheritance and subtyping is the definition of subtyping between session types. Let  $S$  be the set of session types. Define  $\text{unfold}(\mu X.S) = \text{unfold}(\{S/(\mu X.S)\}X)$ , and  $\text{unfold}(S) = S$  for non-recursive session types  $S$ ; contractivity guarantees that this definition terminates.

DEFINITION 1. A relation  $R \subseteq S \times S$  is a session type simulation if  $(S, S') \in R$  implies the following conditions.

1. If  $\text{unfold}(S) = \&\{m_i : S_i\}_{i \in I}$  then  $\text{unfold}(S') = \&\{m_j : S'_j\}_{j \in J}$ ,  $J \subseteq I$  and  $\forall j \in J. (S_j, S'_j) \in R$ .
2. If  $\text{unfold}(S) = \oplus\{E.l_i : S_i\}_{i \in I}$  then  $\text{unfold}(S') = \oplus\{E'.l_j : S'_j\}_{j \in J}$ ,  $\text{constants}(E) \subseteq \text{constants}(E')$  and  $\forall i \in I. (S_i, S'_i) \in R$ .

The subtyping relation on session types is defined by  $S <: S'$  if there exists a session type simulation  $R$  such that  $(S, S') \in R$ .

$$\begin{array}{c}
h; \text{new } C() \longrightarrow h :: o = C[\text{session}(C); \text{fields}(C) = \vec{\text{null}}]; o \quad (\text{R-NEW}) \\
\frac{h(o).f = v}{h; o.f \longrightarrow h + o.f = \text{null}; v} \quad h; o.f = v \longrightarrow h + o.f = v; \text{null} \quad (\text{R-FIELD, R-ASSIGN}) \\
\frac{h(o).f = o' \quad h(o').\text{session} = \&\{m: S, \dots\} \quad m() \{e\} \in \text{methods}(h(o').\text{class})}{h; o.f.m() \longrightarrow h + o.f = \text{null} + o'.\text{session} = S; e\{o'/\text{this}\} \text{ then } f = o'} \quad (\text{R-CALL}) \\
\frac{h(o).f = (E.l_j \text{ then } g = o') \quad h(o').\text{session} = \oplus\{E.l_j: S_j, \dots\} \quad j \in I}{h; \text{switch } (o.f) \{\text{case } l_i: e_i\}_{i \in I} \longrightarrow h + o.f = \text{null} + o.g = o' + o'.\text{session} = S_j; e_j} \quad (\text{R-SWITCH}) \\
\frac{h(o).f = (\text{Bool.F} \text{ then } g = o') \quad h(o').\text{session} = \oplus\{\text{Bool.F}: S, \dots\}}{h; \text{while } (o.f) \{e\} \longrightarrow h + o.f = \text{null} + o.g = o' + o'.\text{session} = S; \text{null}} \quad (\text{R-WHILEF}) \\
\frac{h(o).f = (\text{Bool.T} \text{ then } g = o') \quad h(o').\text{session} = \oplus\{\text{Bool.T}: S, \dots\}}{h; \text{while } (o.f) \{e\} \longrightarrow h + o.f = \text{null} + o.g = o' + o'.\text{session} = S; (e; \text{while } (o.f) \{e\})} \quad (\text{R-WHILET}) \\
h; v; e \longrightarrow h; e \quad \frac{h; e \longrightarrow h'; e'}{h; E[e] \longrightarrow h'; E[e']} \quad (\text{R-SEQ, R-CONTEXT})
\end{array}$$

Figure 11. Reduction rules

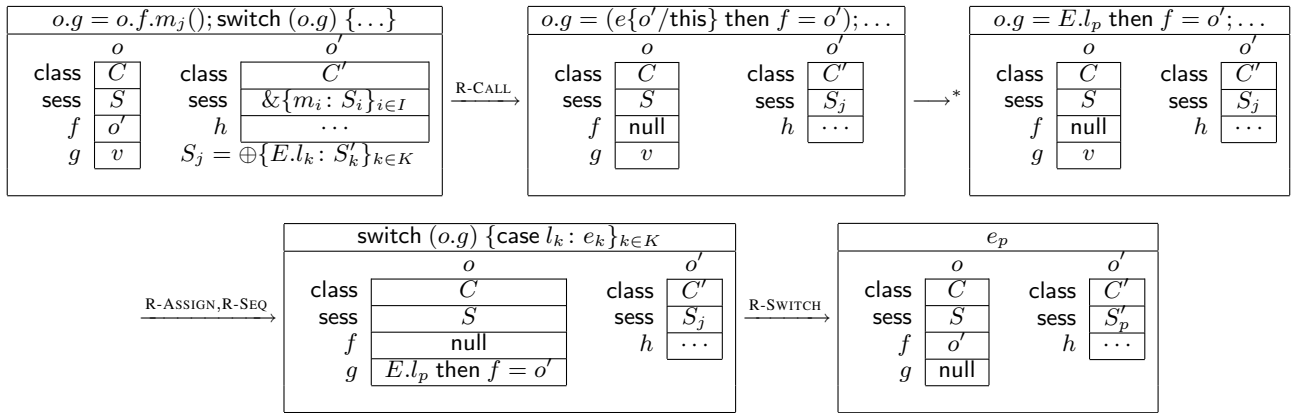


Figure 12. Example of the interplay between method call, switch and the then expression

The direction of subtyping is opposite to that defined in (19), because we make a choice by selecting a method from a session of  $\&$  type instead of by sending a label on a channel of  $\oplus$  type. However, the point is that in both cases, the type allowing a choice to be made has contravariant subtyping in the set of choices. This reversal of the subtyping relation for session types also occurs in (5). Further details, including the proof that subtyping is reflexive and transitive and an algorithm for checking subtyping, can easily be adapted from (19).

Figure 13 defines subtyping between types of our language. The relation is as expected for object types viewed as records of fields, with the addition of subtyping between the session types.

It turns out that both the **requires** and **ensures** types behave covariantly. For **ensures** this is because the type is really part of the result type of the method, describing the implicitly returned **this**. For **requires** it is because the type is the true type of the object on which the method is called.

#### 4.4 Type System

The type system is defined by the rules in Figures 14, 15, and 16. The typing judgement for expressions is  $\vdash \Gamma \triangleright e: T \triangleleft \Gamma'$ . Here  $\Gamma$  and  $\Gamma'$  are the initial and final type environments when typing  $e$ ;  $\Gamma'$  may differ from  $\Gamma$  either because identifiers disappear (due to linearity)

or because their types change (due to their dynamic interfaces). We regard an environment  $\Gamma$  as a function from object references  $o$  to object types  $C[S; \vec{T} \vec{f}]$ . If  $h(o) = C[S; \vec{T} \vec{f}]$  then  $h(o).f_i$  means  $T_i$ .  $\Gamma + o: T$  means disjoint union. If  $\Gamma(o)$  is defined and has field  $f$  (this is an implicit hypothesis) then the notation  $\Gamma + o.f: T$  means the environment obtained by changing the type of field  $f$  in object  $o$  to  $T$ .

First consider Figure 14, which defines typing of expressions. T-NEW types a new object, giving it the initial session type from the class declaration and giving all the fields type Null. T-FIELD types field access, nullifying the field because its value has moved into the expression part of the judgement. T-ASSIGN types field update; the type of the field changes, and the type of the expression is Null, again because of linearity.

T-CALL requires an environment in which method  $o.f.m$  is available. In the signature of  $m$ , the req type must match the type of  $o.f$ , the ens type gives the final type of  $o.f$ , and the result type gives the type of the expression as usual.

T-SWITCH types a switch on  $o.f$ , whose type must have a  $\oplus$  session type and a variant field typing indexed by the same enum  $E$ . In the typing of branch  $e_i$ , the index  $i$  selects from both the session type and the variant type. All branches must have the same

$$\begin{array}{c}
\vdash \Gamma \triangleright \text{null} : \text{Null} \triangleleft \Gamma \quad \vdash \Gamma + o : T \triangleright o : T \triangleleft \Gamma \quad (\text{T-NULL,T-REF}) \\
\frac{l \in \text{constants}(E)}{\vdash \Gamma \triangleright E.l : E \triangleleft \Gamma} \quad \frac{\vdash \Gamma \triangleright e : E \triangleleft \Gamma + o : C[\oplus\{E.l_i : S_i\}_{i \in I}; F_j] \quad j \in I}{\vdash \Gamma \triangleright e : E \triangleleft \Gamma + o : C[\oplus\{E.l_i : S_i\}_{i \in I}; \langle E.l_i : F_i \rangle_{i \in I}]} \quad (\text{T-CONST,T-INJ}) \\
\frac{\vdash \Gamma \triangleright e : T \triangleleft \Gamma' + o : U}{\vdash \Gamma \triangleright e \text{ then } f = o : T \text{ then } U \text{ } f \triangleleft \Gamma'} \quad \vdash \Gamma \triangleright \text{new } C() : C[\text{session}(C); \vec{\text{Null}} \text{ fields}(C)] \triangleleft \Gamma \quad (\text{T-THEN,T-NEW}) \\
\frac{\vdash \Gamma \triangleright o : T \triangleleft \Gamma'}{\vdash \Gamma \triangleright o.f : T.f \triangleleft \Gamma' + o.f : \text{Null}} \quad \frac{\vdash \Gamma \triangleright e : T \triangleleft \Gamma'}{\vdash \Gamma \triangleright o.f = e : \text{Null} \triangleleft \Gamma' + o.f : T} \quad (\text{T-FIELD,T-ASSIGN}) \\
\frac{\Gamma(o).f = T_1 \quad M \in \text{methods}(T_1.\text{class}) \quad \text{req } T_1 \text{ ens } T_2 \text{ for } T \text{ } m() \{e\} <: M}{\vdash \Gamma \triangleright o.f.m() : (T \text{ then } T_2 \text{ } f) \triangleleft \Gamma + o.f : \text{Null}} \quad (\text{T-CALL}) \\
\frac{\Gamma(o).f <: E \text{ then } C[\oplus\{E.l_i : S_i\}_{i \in I}; \langle E.l_i : F_i \rangle_{i \in I}] \text{ } g}{\forall i.l_i \in \text{constants}(E) \quad \vdash \Gamma + o.f : \text{Null} + o.g : C[S_i; F_i] \triangleright e_i : T \triangleleft \Gamma' \quad (\forall i \in I)} \quad (\text{T-SWITCH}) \\
\frac{\Gamma(o).f <: \text{Bool} \text{ then } C[\oplus\{\text{Bool.F} : S_f, \text{Bool.T} : S_t\}; \langle \text{Bool.F} : F_f, \text{Bool.T} : F_t \rangle] \text{ } g \quad \vdash \Gamma + o.f : \text{Null} + o.g : C[S_t; F_t] \triangleright e : \_ \triangleleft \Gamma}{\vdash \Gamma \triangleright \text{while } (o.f) \{e\} : \text{Null} \triangleleft \Gamma + o.f : \text{Null} + o.g : C[S_f; F_f]} \quad (\text{T-WHILE}) \\
\frac{\vdash \Gamma \triangleright e : \_ \triangleleft \Gamma'' \quad \vdash \Gamma'' \triangleright e' : T \triangleleft \Gamma'}{\vdash \Gamma \triangleright e; e' : T \triangleleft \Gamma'} \quad \frac{\vdash \Gamma \triangleright e : T \triangleleft \Gamma' \quad T <: U}{\vdash \Gamma \triangleright e : U \triangleleft \Gamma'} \quad (\text{T-SEQ,T-SUB})
\end{array}$$

**Figure 14.** Typing rules for expressions

$$\begin{array}{c}
\frac{\vdash \text{this} : C[S_j; F] \triangleright e : E \triangleleft \text{this} : C[S_j; \langle E.l_i : F_i \rangle_{i \in I}] \quad j \in I \quad \vdash M_i \quad (\forall i)}{\vdash \text{req } C[\&\{m_i : S_i\}_{i \in I}; F] \text{ ens } C[S_j; \langle E.l_i : F_i \rangle_{i \in I}] \text{ for } E \text{ } m_j() \{e\} \quad \vdash \text{class } \_ \{ \_ ; \_ ; \vec{M} \}} \quad (\text{T-METH,T-CLASS}) \\
\frac{\forall i. \vdash M_i \quad S <: \text{session}(C) \quad M <: M' \quad (\forall M' \in \text{methods}(C'), M \in \vec{M}. \text{name}(M) = \text{name}(M'))}{\vdash \text{class } C \text{ extends } C' \{S; \vec{f}; \vec{M}\}} \quad (\text{T-EXTENDS}) \\
\frac{\vdash \text{enum } E \text{ } L \quad E <: E'}{\vdash \text{enum } E \text{ restricts } E' \_} \quad \frac{\vdash D_i \quad (\forall i)}{\vdash \vec{D}} \quad (\text{T-ENUM,T-RESTRICTS,T-PROG})
\end{array}$$

**Figure 15.** Typing rules for programs

$$\begin{array}{c}
T <: T \quad \frac{T <: T'' \quad T'' <: T'}{T <: T'} \quad (\text{S-ID,S-TRANS}) \quad \vdash \emptyset \triangleright \emptyset \quad \frac{\vdash h \triangleright \Gamma + o : T}{\vdash h \triangleright \Gamma} \quad (\text{T-EMPTY,T-HWEAK}) \\
\frac{\text{enum } E \text{ restricts } E' \text{ } L \in \vec{D} \quad L \subseteq \text{constants}(E')}{E <: E'} \quad (\text{S-ENUM}) \quad \frac{\vdash h \triangleright \Gamma \quad \vdash \Gamma \triangleright \vec{v} : \vec{T} \triangleleft \Gamma' \quad \text{fields}(C) = \vec{f}}{\vdash h :: o = C[S'; \vec{f} = \vec{v}] \triangleright \Gamma' + o : C[S'; \vec{T} \vec{f}]} \quad (\text{T-HADD}) \\
\frac{T <: T' \quad U <: U'}{T \text{ then } U \text{ } f <: T' \text{ then } U' \text{ } f} \quad (\text{S-THEN}) \quad \frac{\vdash \vec{D} \quad \vdash h \triangleright \Gamma \quad \vdash \Gamma \triangleright e : T \triangleleft \Gamma'}{\vdash \Gamma \triangleright h; e : T \triangleleft \Gamma'} \quad (\text{T-STATE}) \\
\frac{\text{class } C \text{ extends } C' \{ \_ ; \_ ; \_ \} \in \vec{D} \quad S <: S' \quad F <: F'}{C[S; F] <: C'[S'; F']} \quad (\text{S-CLASS}) \\
\frac{T_i <: T'_i \quad (\forall i) \quad I \subseteq J \quad E <: E' \quad F_i <: F'_i \quad (\forall i \in I)}{F, \vec{T} \vec{f} <: \vec{T}' \vec{f}'} \quad \frac{\langle E.l_i : F_i \rangle_{i \in I} <: \langle E'.l_j : F'_j \rangle_{j \in J}}{\quad} \quad (\text{S-FIELDS,S-VARIANT}) \\
\frac{T <: T' \quad U <: U' \quad V <: V'}{\text{req } T \text{ ens } U \text{ for } V \text{ } m() \{ \_ \} <: \text{req } T' \text{ ens } U' \text{ for } V' \text{ } m() \{ \_ \}} \quad (\text{S-METHOD})
\end{array}$$

**Figure 16.** Typing rules for heaps and states

**Figure 13.** Subtyping rules for types and method signatures

final environment  $\Gamma'$ , so that it is a consistent final environment for the switch expression.

T-WHILE is based on the idea that a while loop is a switch with two branches. The terminating branch is effectively null and does not appear as a hypothesis. The looping branch, which is the expression  $e$ , must be typable and must preserve the environment, so that the loop makes sense. T-SEQ and T-SUB are standard.

Now consider Figure 15. The most interesting rule is T-METH, which checks that a method body has the effect specified by the req and ens declarations.

Figure 16 defines rules for typing heaps and states (runtime configurations). The typing of a heap,  $\vdash h \triangleright \Gamma$ , means that  $\Gamma$  gives types to the usable objects in  $h$ . Because of linearity,  $\Gamma$  only contains types for top-level objects, i.e. those that are not stored in fields of other objects. Weakening of the heap typing (T-HWEAK)

is needed in order to prove type preservation, because assignment can discard an object.

#### 4.5 Results

By standard techniques (41) adapted to typing judgements with initial and final environments (20) we can prove the usual results: Type Preservation and Type Safety.

**THEOREM 1 (Type Preservation).** *If  $\vdash \Gamma \triangleright h; e : T \triangleleft \Gamma'$  and  $h; e \longrightarrow h'; e'$  then there exists  $\Gamma''$  such that  $\vdash \Gamma'' \triangleright h'; e' : T \triangleleft \Gamma'$ .*

**DEFINITION 2 (Stuck States).** *If  $e$  is neither a value nor a variable and  $\neg \exists h', e'. (h; e \longrightarrow h'; e')$  then the state  $h; e$  is stuck.*

A state is stuck if either the left-hand side of the reduction rule for its syntactic category is undefined, or if some premise does not hold. To prove that the language is type safe, we prove the following by induction on the structure of  $e$ .

**THEOREM 2 (Type Safety).** *If  $\vdash \Gamma \triangleright h; e : T \triangleleft \Gamma'$  then  $h; e$  is not stuck.*

## 5. Typechecking Algorithm

Figure 17 defines a typechecking algorithm. It is used in two steps. First, for each class  $C$  with declared session type  $S$ ,  $\mathcal{P}(C, S, \emptyset)$  is called. This returns annotations for the methods of  $C$ , in the form  $\text{req } C[\&\{m_i : S_i\}_{i \in I}] \text{ ens } C[\Delta(S_i)] \text{ for } T m_i \{e\}$ . Algorithm  $\mathcal{P}$  is very simple, and just translates the session type of a class into explicit pre- and post-conditions for its methods. A particular method can receive several different annotations, giving a form of overloading which is useful in the example of Figure 4, allowing close to be called in three different states. In this case, the req type should be used to disambiguate methods calls.

The second step is to call  $\mathcal{A}(C, F, S, \emptyset)$  for each class  $C$ , where  $S$  is the declared session type of  $C$  and  $F$  is the initial field typing (with all fields having type Null) of  $C$ . Algorithm  $\mathcal{A}$  has two purposes. (1) It calls algorithm  $\mathcal{B}$  to type-check the method bodies of  $C$ , in the order corresponding to  $S$ . While typechecking, the annotations calculated by algorithm  $\mathcal{P}$  are used to check the effect of method calls. (2) It calculates another set of annotations for the methods of  $C$ , in the form  $\text{req } C[\&\{m_i : S_i\}_{i \in I}; F] \text{ ens } C[S_j; F_j] \text{ for } T m_j() \{e\}$ . These are used in the proof of type safety, to show that a typable program in the top-level language yields a typable program in the runtime language.

The definition of  $\mathcal{B}$  follows the typing rules (Figure 14) except for one point: T-INJ means that the rules are not syntax-directed. To compensate, clause  $E.l$  of  $\mathcal{B}$  produces a *partial* variant field typing with an incomplete set of labels, and clause switch uses the  $\uplus$  operator to combine partial variants and check for consistency.

The various “where” and “if” clauses should be interpreted as conditions for the functions to be defined; cases in which the functions are undefined should be interpreted as typing errors.

For example, applying algorithm  $\mathcal{P}$  to class FileReaderToEnd in Figure 1 produces the following annotated methods.

```

req FileReaderToEnd [Init]
ens FileReaderToEnd [ $\oplus$ {Res.OK: Open, ...}]
Res open ()

req FileReaderToEnd [Open]
ens FileReaderToEnd [ $\oplus$ {Bool.True: Close, ...}]
Bool eof ()

req FileReaderToEnd [Read]
ens FileReaderToEnd [Open]

```

#### Algorithm $\mathcal{P}$

$$\begin{aligned} \mathcal{P}(C, \&\{m_i : S_i\}_{i \in I}, \Delta) = & \\ & \{ \text{req } C[\&\{m_i : S_i\}_{i \in I}] \text{ ens } C[\Delta(S_i)] \text{ for } T m_i \{e\} \\ & \mid T m_i() \{e\} \in \text{methods}(C) \} \cup \bigcup_{j \in I} \mathcal{P}(C, S_j, \Delta) \\ \mathcal{P}(C, \oplus\{E.l_i : S_i\}_{i \in I}, \Delta) = & \bigcup_{i \in I} \mathcal{P}(C, S_i, \Delta) \\ \mathcal{P}(C, \mu X.S, \Delta) = & \mathcal{P}(C, S, \Delta \cup \{X \mapsto \mu X.S\}) \end{aligned}$$

#### Algorithm $\mathcal{A}$

$$\begin{aligned} \mathcal{A}(C, F, \&\{m_i : S_i\}_{i \in I}, \Delta) = & \\ \bigcup_{j \in I} \{ \text{req } C[\&\{m_i : S_i\}_{i \in I}; F] \text{ ens } C[S_j; F_j] \text{ for } T m_j() \{e\} \} & \\ \cup \mathcal{A}(C, F_j, S_j, \Delta) \mid T m_j() \{e\} \in \text{methods}(C), & \\ (F_j, T) = \mathcal{B}(F, e) \} & \\ \mathcal{A}(C, F, \oplus\{E.l_i : S_i\}_{i \in I}, \Delta) = & \bigcup_{j \in I} \mathcal{A}(C, F, S_j, \Delta) \\ \mathcal{A}(C, F, \mu X.S, \Delta) = & \mathcal{A}(C, F, S, \Delta \cup \{X \mapsto F\}) \\ \mathcal{A}(C, F, X, \Delta) = & \text{if } \Delta(X) = F \text{ then } \emptyset \end{aligned}$$

#### Algorithm $\mathcal{B}$

$$\begin{aligned} \mathcal{B}(F, \text{null}) = & (F, \text{Null}) \\ \mathcal{B}(F, E.l) = & ((E.l : F), E) \\ \mathcal{B}(F, \text{new } C()) = & (F, C[\text{session}(C)]) \\ \mathcal{B}(F, \text{this}.f) = & (F + f : \text{Null}, F(f)) \\ \mathcal{B}(F, \text{this}.f = e) = & (F' + \text{this}.f : T, \text{Null}) \\ & \text{where } (F', T) = \mathcal{B}(F, e) \\ \mathcal{B}(F, \text{this}.f.m()) = & (F + f : \text{Null}, T \text{ then } C[S'] f) \\ & \text{where } F(f) = C[S] \text{ and} \\ & \text{req } C[S] \text{ ens } C[S'] \text{ for } T m_j() \{e\} \in \text{methods}(\mathcal{P}(C)) \\ \mathcal{B}(F, \text{switch } (\text{this}.f) \{ \text{case } l_i : e_i \}_{i \in I} \}) = & (\biguplus_{i \in I} F'_i, T) \\ & \text{where } F(f) = (T \text{ then } C[\oplus\{E.l_i : S_i\}_{i \in I}] g) \text{ and} \\ & \forall i \in I. ((F'_i, T) = \mathcal{B}(F + f : \text{Null} + g : C[S_i], e_i)) \\ \mathcal{B}(F, \text{while } (\text{this}.f) \{e\}) = & (F_f + f : \text{Null} + g : C[S_f], \text{Null}) \\ & \text{where } F(f) = (\text{Bool then } C[\oplus\{\text{Bool.F} : S_f, \text{Bool.T} : S_t\}]) \\ & \text{and } (F, T) = \mathcal{B}(F_t, e) \\ \mathcal{B}(F, e; e') = & \mathcal{B}(\vec{m}, F', e') \text{ where } (F', -) = \mathcal{B}(\vec{C}, F, e) \end{aligned}$$

#### Combining partial variants

$$\begin{aligned} F \uplus F = F & \text{ if } F \text{ is not a variant typing} \\ \langle E.l_i : F_i \rangle_{i \in I} \uplus \langle E.m_j : G_j \rangle_{j \in J} = \langle E.l_k : H_k \rangle_{k \in I \cup J} & \\ \text{where } \forall i \in I. (H_i = F_i), \forall j \in J. (H_j = G_j), & \\ \text{and whenever } l_i = l_j \text{ we have } F_i = G_j & \end{aligned}$$

Figure 17. Algorithm

```
String read ()
```

```

req FileReaderToEnd [Close]
ens FileReaderToEnd [end]
Null close ()

```

These annotations are used when algorithm  $\mathcal{A}$  is applied to class FileReader in Figure 2 producing the following annotated methods. These annotations include information about the field typing within FileReader, but no information about fields within FileReaderToEnd, for which we do not have the source code.

```

req FileReader [Init; Null f]
ens FileReader [Read; FileReaderToEnd [Init] f]
Null init () {...}

req FileReader [Read; FileReaderToEnd [Init] f]
ens FileReader [Final; FileReaderToEnd [end] f]
Null read () {...}

req FileReader [Final; FileReaderToEnd [end] f]
ens FileReader [Final; FileReaderToEnd [end] f]

```



String toString() {...}

**THEOREM 3.** Let  $\vec{D}$  be a program, i.e. a sequence of declarations.

1. For every class  $C \{S; \vec{f}; \vec{M}\}$  or class  $C$  extends  $C' \{S; \vec{f}; \vec{M}\}$  in  $\vec{D}$ ,  $\mathcal{P}(C, S, \emptyset)$  and  $\mathcal{A}(C, \text{Null}\vec{f}, S, \emptyset)$  terminate.
2. For every class  $C \{S; \vec{f}; \vec{M}\}$  or class  $C$  extends  $C' \{S; \vec{f}; \vec{M}\}$  in  $\vec{D}$ , replace  $\vec{M}$  by  $\mathcal{A}(C, \text{Null}\vec{f}, S, \emptyset)$ , and let  $\vec{D}'$  be the resulting declarations. Then  $\vdash \vec{D}'$ .

The session type of a class has two interpretations. The first is as a limit on the allowed sequences of method calls, a kind of *safety* property, and this is always guaranteed by our type safety result. The second interpretation is that every sequence of method calls in the session type should be realizable in a typable program. Given a class definition  $C$  in the internal syntax, with explicit req and ens annotations, construct an expression  $e_C$  as follows:

1. View the session type of  $C$  as a (possibly infinite) tree, with branching at  $\&$  and  $+$  nodes.
2. Make it into a finite tree by replacing some  $\&$  nodes by end.
3. For each  $\&$  node, remove all except one of the branches; call the resulting tree  $T$ .
4.  $e_C$  constructs an object of class  $C$  and contains a sequence of method calls and switch statements corresponding to the structure of  $T$ .

Typability of  $C$  in the internal system does not guarantee that  $e_C$  is typable, because it is possible for the req and ens clauses to contain spurious constraints such that the ens of one method does not match the req of the next method in the session type. But the typechecking algorithm, applied to a program in the programmer's syntax, produces definitions such that every  $e_C$  is typable.

## 6. Related Work

**Cyclone, Vault, CQual, Fugue, Sing#.** *Cyclone* (21; 22), *Vault* (9; 16), and *CQual* (18) are systems based on the C programming language that allow protocols to be statically enforced by a compiler. *Cyclone* adds many benefits to C, but its support for protocols is limited to enforcing locking of resources. Between acquiring and releasing a lock, there are no restrictions on how a thread may use a resource. In contrast, our system uses types both to enforce locking of objects (via linearity) and to enforce the correct sequence of method calls.

*Vault* is much closer to our system, allowing abstract states to be defined for resources, with pre- and post-conditions for each operation, and checking statically that operations occur in the correct sequence. It uses linear types to control aliasing, and uses the *adoption and focus* mechanism (16) to re-introduce aliasing in limited situations. *Fugue* (17; 10) extends similar ideas to an object-oriented language, and uses explicit pre- and post-conditions that are somewhat similar to our req/ens annotations. *CQual* expects users to annotate programs with type qualifiers; its type system, simpler and less expressive than the above, provides for type inference. The main novelty of our work is the use of the session type of a class as a global specification, to introduce a dependency between the result of a method and the subsequent abstract state of the object, and to characterize the subtyping relation.

*Sing#* (15) is an extension of C# which has been used to implement Singularity, an operating system based on message-passing. It incorporates session types to specify protocols for communication channels, and introduces *contracts* which are analogous to our req and ens clauses. Like our system, *Sing#* is object-oriented, but it uses dynamic interfaces only in relation to communication

channels, and not with classes in general. Although inheritance of contracts is mentioned very briefly in (15), details do not seem to have been published; in our system the use of the session type as a global specification is significant in describing inheritance and subtyping. A technical point is that *Sing#* uses a single construct switch receive to combine receiving an enumeration value and doing a case-analysis, whereas our system allows a switch on an enumeration value to be separated from the method call that produces it.

**Resource Usage Analysis.** Igarashi and Kobayashi (28) define a general resource usage analysis problem for an extended  $\lambda$ -calculus, including a type inference system, that statically checks the order of resource usage. Albeit quite expressive, their system does not seem to allow expressing branching on method results.

**Non-Uniform Concurrent Objects / Active Objects.** Another related line of research was started by Nierstrasz (31), aimed at describing the behaviour of non-uniform *active* objects in concurrent systems, whose behaviour (including the set of available methods) may change dynamically. He defined subtyping for active objects, but did not formally define a language semantics or a type system. The topic has been continued by several authors including Ravara *et al.* (38; 39), Puntigam and Peter (35; 36), and Caires (4). Damiani, Giachino, et.al. (8), propose a Java-like concurrent language incorporating inheritance and subtyping and equipped with a type-and-effect system, where method availability is made dependent on the state of objects. Our language is sequential, but we go beyond all existing work on non-uniform objects referred above by describing client branching on method results. Caires' work has finer-grained control of resources.

**Unique Ownership of Objects.** Our system treats objects with dynamic interfaces linearly in a straightforward way to prevent the creation of aliases and thus avoid the problem of maintaining consistency between different views of an object. This idea goes back at least as far as (2) and has been applied many times. However, linearity causes problems of its own: linear objects cannot be stored in shared data structures, and this tends to restrict expressivity. There is a large literature on less extreme techniques for static control of aliasing: Hogg's *Islands* (23), Almeida's *balloon types* (1), Noble *et al.*'s *flexible alias protection* (32), Clarke *et al.*'s *ownership types* (7), Föhndrich and DeLine's *adoption and focus* (16), Östlund *et al.*'s *Joe<sub>3</sub>* (33) among others.

We have developed our type system for dynamic interfaces, and the theory of inheritance, using the simplest possible approach to alias control. We hope that it will be possible to integrate a more sophisticated alias analysis, such as those mentioned above, into our system. The key property we need is that when changing the type of an object (by calling a method on it or by performing a switch or a while on an enumeration constant returned from a method call) there must be a unique reference to it.

**Session Types for Object-Oriented Languages.** Several recent papers by Dezani-Ciancaglini, Yoshida *et al.* (11; 12; 13; 14; 26) have combined session types, as specifications of protocols on communication channels, with the object-oriented paradigm. This work focuses on communication channels and does not address the more general question of dynamic interfaces. Their approach is also different from that of *Sing#*. Instead of defining a contract for a channel, with methods for sending and receiving, a channel is created and used entirely within a single method. Bounded polymorphism, which includes subtyping, is studied in (14) but it seems orthogonal to the question of inheritance and subclasses.

**Analysis of Concurrent Systems using Pi-Calculus.** Some work on static analysis of concurrent systems expressed in pi-calculus is also relevant, in the sense that it addresses the question (among others) of whether attempted uses of a resource are consistent with its state. Kobayashi *et al.* have developed a generic framework (27)

including a verification tool (29) in which to define type systems for analyzing various behavioural properties including sequences of resource uses (30). In some of this work, types are themselves abstract processes, and therefore in some situations resemble our session types. Rajamani *et al.*'s *Behave* (6; 37) uses CCS to describe properties of  $\pi$ -calculus programs, and verify the validity of temporal formulae via a combination of type-checking and model-checking techniques, thereby going beyond static analysis. However, all of this work is based on  $\pi$ -calculus; it is not obvious how to use these results to directly tackle the problem of analysis of object-oriented programs.

## 7. Future Work

There are several topics for future work.

**Shared classes.** In the present system, all classes are linear. It is straightforward to add shared classes, whose objects do not have to be uniquely referenced. The behaviour of shared objects is largely orthogonal to that of linear objects, except for the condition that a shared object's fields cannot contain linear objects.

**More general methods.** It should not be difficult to allow methods to return arbitrary values, not just enumerated values, meaning that session types would not have to alternate between  $\&$  and  $\oplus$ . Supporting method parameters is very straightforward.

**More flexible control of aliasing.** The mechanism for controlling aliasing should be orthogonal to the theory of how operations affect uniquely-referenced objects. We intend to adapt existing work to relax our strictly linear control and obtain a more flexible language.

**Self calls.** For technical reasons, the present version of our work does not allow self calls (`this.m(...)`). One solution could be to distinguish between *private* and *public* session types, related by hiding some methods; self calls would be allowed for private methods.

**Complete use of sessions.** Some systems based on session types guarantee that sessions are completely used, finishing in state end. Our system does not have this property, but to achieve it we only need to change the rules for assignment so that an incompletely-used object cannot be discarded. In a practical language this condition could be specified independently for each object.

**Java-style interfaces.** If class  $C$  implements interface  $I$  then we should have  $\text{session}(C) <: \text{session}(I)$ , interpreting the interface as a specification of minimum method availability.

**Concurrency.** We believe that by extending our system to a concurrent language with channel-based communication, we can unify dynamic interfaces, the *Sing#* (15) version of session types, and other work on object-oriented session types (12; 13; 14; 26). It will also be interesting to look at RMI for dynamic interfaces.

**Specifications involving several objects.** Multi-party session types (3; 24) allow to describe specifications that involve more than two objects. The introduction of such types would allow disciplining the usage of more sophisticated patterns of object usage.

## 8. Conclusions

We have defined a core class-based object-oriented language with a static type system which is able to check correctness of method calls in a setting in which method availability depends on an object's state. We have proved the correctness of the type system with respect to a formal operational semantics. Key features of the language are: firstly, allowing a dependency between the result of a method and the subsequent abstract state of its object, thereby forcing the client of such a method to branch accordingly; secondly, the use of a session type as a global representation of the abstract states of an object, enabling us to define rules for inheritance in terms of a subtyping relation on session types.

## Acknowledgments

Simon Gay was partially supported by the Security and Quantum Information Group, Instituto de Telecomunicações, Portugal, and by the UK EPSRC grants "Engineering Foundations of Web Services: Theories and Tool Support" (EP/E065708/1) and "Behavioural Types for Object-Oriented Languages" (EP/F037368/1). He thanks the University of Glasgow for the sabbatical leave during which part of this research was done.

António Ravara and Vasco T. Vasconcelos were partially supported by FEDER, the EU IST proactive initiative FET-Global Computing (project Sensoria, IST-2005-16004), and by the Portuguese FCT (via SFRH/BSAB/757/2007, and project Space-Time-Types, POSC/EIA/55582/2004).

António Ravara was also partially supported by the UK EPSRC grant "Behavioural Types for Object-Oriented Languages" (EP/F037368/1).

## References

- [1] P. S. Almeida. Balloon types: Controlling sharing of state in data types. *ECOOP, Springer LNCS*, 1241:32–59, 1997.
- [2] H. G. Baker. 'use-once' variables and linear objects — storage management, reflection and multi-threading. *ACM SIGPLAN Notices*, 30(1), 1995.
- [3] E. Bonelli and A. Compagnoni. Multipoint Session Types for a Distributed Calculus. In G. Barthe and C. Fournet, editors, *Proceedings of the Third International Symposium on Trustworthy Global Computing*, volume 4912 of *Lecture Notes in Computer Science*, pages 240–256. Springer-Verlag, 2007.
- [4] L. Caires. Spatial-behavioral types, distributed services, and resources. *TGC, Springer LNCS*, 4661:98–115, 2006.
- [5] M. Carbone, K. Honda, and N. Yoshida. Structured global programming for communication behaviour. *ESOP, Springer LNCS*, 4421:2–17, 2007.
- [6] S. Chaki, S. K. Rajamani, and J. Rehof. Types as models: model checking message-passing programs. In *POPL*, pages 45–57. ACM Press, 2002.
- [7] D. G. Clarke, J. M. Potter, and J. Noble. Ownership types for flexible alias protection. In *OOPSLA*, volume 33(10) of *SIGPLAN Notices*. ACM Press, 1998.
- [8] F. Damiani, E. Giachino, P. Giannini, and S. Drossopoulou. A type safe state abstraction for coordination in Java-like languages. *Acta Informatica*, October 2008.
- [9] R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI*, volume 36(5) of *SIGPLAN Notices*, pages 59–69. ACM Press, 2001.
- [10] R. DeLine and M. Fähndrich. The Fugue protocol checker: is your software Baroque? Technical Report MSR-TR-2004-07, Microsoft Research, 2004.
- [11] M. Dezani-Ciancaglini, S. Drossopoulou, E. Giachino, and N. Yoshida. Bounded session types for object-oriented languages. *FMCO, Springer LNCS*, 4709:207–245, 2007.
- [12] M. Dezani-Ciancaglini, D. Mostrous, N. Yoshida, and S. Drossopoulou. Session types for object-oriented languages. *ECOOP, Springer LNCS*, 4067:328–352, 2006.
- [13] M. Dezani-Ciancaglini, N. Yoshida, A. Ahern, and S. Drossopoulou. A distributed object-oriented language with session types. *TGC, Springer LNCS*, 3705:299–318, 2005.
- [14] S. Drossopoulou, M. Dezani-Ciancaglini, and M. Coppo. Amalgamating the session types and the object oriented programming paradigms. Manuscript, 2007.
- [15] M. Fähndrich, M. Aiken, C. Hawblitzel, O. Hodson, G. Hunt, J. R. Larus, and S. Levi. Language support for fast and reliable message-based communication in Singularity OS. In *EuroSys*. ACM Press, 2006.

- [16] M. Fähndrich and R. DeLine. Adoption and focus: practical linear types for imperative programming. In *PLDI*, volume 37(5) of *SIGPLAN Notices*, pages 13–24. ACM Press, 2002.
- [17] M. Fähndrich and R. DeLine. Typstates for objects. *ESOP, Springer LNCS*, 3086:465–490, 2004.
- [18] J. S. Foster, T. Terauchi, and A. Aiken. Flow-sensitive type qualifiers. In *PLDI '02: Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation*, pages 1–12, New York, NY, USA, 2002. ACM.
- [19] S. J. Gay and M. J. Hole. Subtyping for session types in the pi calculus. *Acta Informatica*, 42(2/3):191–225, 2005.
- [20] S. J. Gay, A. Ravara, and V. T. Vasconcelos. Session types for inter-process communication. Technical Report TR-2003-133, Department of Computing Science, University of Glasgow, March 2003.
- [21] D. Grossman. Type-safe multithreading in Cyclone. In *TLDI*, volume 38(3) of *SIGPLAN Notices*, pages 13–25. ACM Press, 2003.
- [22] D. Grossman, G. Morrisett, T. Jim, M. Hicks, Y. Wang, and J. Cheney. Region-based memory management in Cyclone. In *PLDI*, volume 37(5) of *SIGPLAN Notices*, pages 282–293. ACM Press, 2002.
- [23] J. Hogg. Islands: aliasing protection in object-oriented languages. In *OOPSLA*, volume 26(11) of *SIGPLAN Notices*, pages 271–285. ACM Press, 1991.
- [24] K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In G. C. Necula and P. Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008*, pages 273–284. ACM, 2008.
- [25] K. Honda, V. Vasconcelos, and M. Kubo. Language primitives and type discipline for structured communication-based programming. *ESOP, Springer LNCS*, 1381:122–138, 1998.
- [26] R. Hu, N. Yoshida, and K. Honda. Session-based distributed programming in Java. *ECOOP, Springer LNCS*, 5142:516–541, 2008.
- [27] A. Igarashi and N. Kobayashi. A generic type system for the pi-calculus. *Theoretical Computer Science*, 311(1-3):121–163, 2004.
- [28] A. Igarashi and N. Kobayashi. Resource usage analysis. *ACM Trans. on Programming Languages and Systems*, 27(2):264–313, 2005.
- [29] N. Kobayashi. Type-based information flow analysis for the pi-calculus. *Acta Informatica*, 42(4–5):291–347, 2005.
- [30] N. Kobayashi, K. Suenaga, and L. Wischik. Resource usage analysis for the  $\pi$ -calculus. *Logical Methods in Computer Science*, 2(3:4):1–42, 2006.
- [31] O. Nierstrasz. Regular types for active objects. In *Object-Oriented Software Composition*, pages 99–121. Prentice Hall, 1995.
- [32] J. Noble, J. Vitek, and J. Potter. Flexible alias protection. *ECOOP, Springer LNCS*, 1445:158–185, 1998.
- [33] J. Östlund, T. Wrigstad, D. Clarke, and B. Åkerblom. Ownership, uniqueness and immutability. In *International Workshop on Aliasing, Confinement and Ownership in object-oriented programming (IWACO)*, 2007.
- [34] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [35] F. Puntigam. State inference for dynamically changing interfaces. *Computer Languages*, 27:163–202, 2002.
- [36] F. Puntigam and C. Peter. Types for active objects with static deadlock prevention. *Fundamenta Informaticae*, 49:1–27, 2001.
- [37] S. K. Rajamani and J. Rehof. A behavioral module system for the pi-calculus. *SAS, Springer LNCS*, 2126:375–394, 2001.
- [38] A. Ravara and L. Lopes. Programming and implementation issues in non-uniform TyCO. Technical report, Department of Computer Science, Faculty of Sciences, University of Porto, Portugal, 1999.
- [39] A. Ravara and V. T. Vasconcelos. Typing non-uniform concurrent objects. *CONCUR, Springer LNCS*, 1877:474–488, 2000.
- [40] K. Takeuchi, K. Honda, and M. Kubo. An interaction-based language and its typing system. *PARLE, Springer LNCS*, 817, 1994.
- [41] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.