

Behavioural types for memory safety in Mungo

António Ravara

With M. Bravetti (IT), A. Francalanza (MT), H. Hüttl (DK)

NOVA LINCS and Departamento de Informática
Faculdade de Ciências e Tecnologia, Universidade NOVA de Lisboa
Work partially supported by RISE and COST (EU), and EPSRC (UK)

December 19, 2019



Research partly supported by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement No 778233.



Aim of type systems: avoid programs to go wrong

Well-typed programs do not go wrong

R. Milner. A Theory of Type Polymorphism in Programming.
Journal of Computer and System Sciences, 1978.

Typical wrong situations in OO prevented by type systems

- ▶ data-mismatch; method-not-understood
(mostly) prevent by classical type systems
- ▶ flow-mismatch, protocol violations
prevent by behavioural type systems

Behavioural types in Java-like languages

Associate with a class a description of an object's behaviour.

Class usages

- ▶ Declare the admissible sequences of method calls.
- ▶ Ensure at compile time that sequences of method calls follow the order declared by usages.
- ▶ Objects are linear to prevent interferences unexpectedly changing their state.

Some languages

- ▶ Campos and Vasconcelos MOOL – gloss.di.fc.ul.pt/mool
- ▶ Gay et al. Mungo – www.dcs.gla.ac.uk/research/mungo/
- ▶ Aldrich et al. Plaid – www.cs.cmu.edu/~aldrich/plaid/
- ▶ Hu et al. SJ – www.doc.ic.ac.uk/~rhu/sessionj.html

Class usage types, by example: a file

```

Class File {
1
2
3 // usage = {open; <rec X.
4 //   {neof; <true:{read;X}, false:{close;end}}},end}}
5
6   bool open() {...}
7
8   bool neof() {...}
9
10  String read() {...}
11
12  void close() {...}
13
14 }

```

Class usages, by example: a “client” of the file

```

Class FileReader { 1
// usage = {init;{fetchData;{readData;end}}}} 2
  File f; String s; 3
  void init() { 4
    f = new File(); 5
    s = ""; 6
  } 7
  void fetchData () { 8
    if (f.open()) { 9
      while (f.neof()) 10
        s ⊕= f.read(); 11
      f.close(); 12
    } 13
  String readData() { s; } 14
} 15

```

What can still go wrong?

Null de-referencing

Forget to initialise a field.

```

class FileReader {
    // usage = {open; <rec X.
    // {neof; <true:{read;X}, false:{close;end}}},end}}
    File f;

    void init() {
        f.open();
    }
    ...
}

```

1
2
3
4
5
6
7
8
9
10
11
12
13

What can still go wrong?

Dangling pointers

- ▶ Re-initialise fields with linear objects.
- ▶ Pass a linear object as an argument of a method call and do not return it even if it did not complete its protocol.

```

class FileReader {
  ...
  void fetchData () {
    if (f.open()) {
      f = new File();
      while (f.neof())
        s ⊕= f.read();
      f.close();
    }
    ...
  }
}

```

1
2
3
4
5
6
7
8
9
10
11

Aim of type systems: avoid programs to go wrong

- ▶ Forgetting to initialise a field before calling a method on it results in a null-pointer dereferencing, even though you follow the protocol (fidelity).
- ▶ If one loses reference to an object in the middle of a protocol, it may not be completed.
- ▶ OO type systems (even behavioural ones) DO NOT statically catch this situations.
- ▶ Languages (extensions) dealing with the issues require annotations.

What do we want to ensure?

Less wrong well-typed programs, by

- ▶ performing null-pointer analysis, and
- ▶ better linearity control to ensure protocol completion.

Aim of type systems: avoid programs to go wrong

- ▶ Forgetting to initialise a field before calling a method on it results in a null-pointer dereferencing, even though you follow the protocol (fidelity).
- ▶ If one loses reference to an object in the middle of a protocol, it may not be completed.
- ▶ OO type systems (even behavioural ones) DO NOT statically catch this situations.
- ▶ Languages (extensions) dealing with the issues require annotations.

What do we want to ensure?

Less wrong well-typed programs, by

- ▶ performing null-pointer analysis, and
- ▶ better linearity control to ensure protocol completion.

The Mungo language

(usage types) $u ::= \{m_i; w_i \mid i \in I\} \mid \mu X.u$

(choice usage types) $w ::= u \mid \langle l_i; u_i \mid i \in I \rangle \mid X$

(basic types) $b ::= \text{void} \mid \text{bool} \mid L$

(types) $t ::= b \mid \perp \mid C[u]$

(enum and class declarations) $D ::= \text{enum } L \{\tilde{l}\} \mid \text{class } C \{u; \vec{F}; \vec{M}\}$

(field and method declarations) $F ::= t f \quad M ::= t m(t x)\{e\}$

(values) $v ::= \text{unit} \mid \text{true} \mid \text{false} \mid l \mid \text{null}$

(references) $r ::= x \mid f$

(expressions) $e ::= v \mid r \mid \text{new } C \mid f = e \mid r.m(e) \mid e; e$
 $\mid \text{if } (e) \{e\} \text{ else } \{e\} \mid \text{switch } (r.m(e)) \{\vec{l} : \vec{e}\}_{l \in L}$
 $\mid \lambda : e \mid \text{continue } \lambda$

Key features of our type-checking system

Null-pointer analysis

A dedicated type \perp inhabited only by *null* that inhabits no other type.

BTW, the fact that *null* is a value of every type is the problem now with type-safety in Java and Scala!

Protocol completion ensures memory safety

We took a step further, not allowing objects to be dropped.

Efficient type-checking

Following the typestate may lead to check the same method as many times as it is mentioned.

Mungo's type-system should check methods only once by inferring the usage of fields and parameters.

Key features of our type-checking system

Null-pointer analysis

A dedicated type \perp inhabited only by *null* that inhabits no other type.

BTW, the fact that *null* is a value of every type is the problem now with type-safety in Java and Scala!

Protocol completion ensures memory safety

We took a step further, not allowing objects to be dropped.

Efficient type-checking

Following the typestate may lead to check the same method as many times as it is mentioned.

Mungo's type-system should check methods only once by inferring the usage of fields and parameters.

Key features of our type-checking system

Null-pointer analysis

A dedicated type \perp inhabited only by *null* that inhabits no other type.

BTW, the fact that *null* is a value of every type is the problem now with type-safety in Java and Scala!

Protocol completion ensures memory safety

We took a step further, not allowing objects to be dropped.

Efficient type-checking

Following the typestate may lead to check the same method as many times as it is mentioned.

Mungo's type-system should check methods only once by inferring the usage of fields and parameters.

Guarantees (I): data-safety and no null de-referencing

Errors to avoid

$$\frac{s(x) = \text{null}}{\vdash_{\vec{D}} \langle h, (o, s) : \text{env}_S, x.m(v) \rangle \longrightarrow_{\text{err}}}$$

$$\frac{h(o).f = \text{null}}{\vdash_{\vec{D}} \langle h, (o, s) : \text{env}_S, f.m(v) \rangle \longrightarrow_{\text{err}}}$$

How to avoid them

$$\frac{\text{lin}(t)}{\Lambda; \Delta : (o, [x \mapsto t]) \vdash_{\vec{D}} x : t \triangleright \Lambda; \Delta : (o, [x \mapsto \perp])}$$

$$\frac{\text{lin}(t)}{\Lambda, o.f \mapsto t; \Delta : (o, S) \vdash_{\vec{D}} f : t \triangleright \Lambda, o.f \mapsto \perp; \Delta : (o, S)}$$

$$\text{TCALL} \frac{\Lambda \vdash_{\vec{D}}^o e : t \triangleright \Lambda', o.r : C[U] \quad U \xrightarrow{m} W \quad t' m(tx) \{e'\} \in C.\text{methods}}{\Lambda \vdash_{\vec{D}}^o r.m(e) : t' \triangleright \Lambda', o.r : C[W]}$$

Guarantees (II): protocol completion

Protocol fidelity, *i.e.*, objects will *not* not follow the protocols.

Will objects actually follow the protocols?

Avoid losing reference to objects with incomplete protocols.

$$\frac{\text{lin}(v', h) \wedge v \neq v'}{\vdash_{\vec{D}} \langle h, (o, [x \mapsto v']) \rangle : \text{env}_S, \text{return}\{v\} \rangle \longrightarrow_{\text{err}}}$$

$$\frac{v \notin \{l_i \mid i \in I\} \quad h(o).\text{usage} = \langle l_i; u_i \mid i \in I \rangle}{\vdash_{\vec{D}} \langle h, (o, s) \rangle : \text{env}_S, \text{return}\{v\} \rangle \longrightarrow_{\text{err}}}$$

$$\frac{\text{lin}(v, h)}{\vdash_{\vec{D}} \langle h, \text{env}_S, v; e \rangle \longrightarrow_{\text{err}}}$$

$$\frac{\Lambda; \Delta \vdash_{\vec{D}} e : t \triangleright \Lambda'; \Delta' \quad \Delta' = \Delta'' : (o', S') \quad \text{terminated}(S')}{\Lambda; \Delta : (o, S) \vdash_{\vec{D}} \text{return}\{e\} : t \triangleright \Lambda'; \Delta' : (o, S)}$$

$$\frac{\Lambda; \Delta \vdash_{\vec{D}} e : t \triangleright \Lambda''; \Delta'' \quad \neg \text{lin}(t) \quad \Lambda''; \Delta'' \vdash_{\vec{D}} e' : t' \triangleright \Lambda'; \Delta'}{\Lambda; \Delta \vdash_{\vec{R}} e; e' : t' \triangleright \Lambda'; \Delta'}$$

The real challenge: infer usages

- ▶ Specify object protocols when writing OO code is not common practice.
- ▶ Is it reasonable to ask programmers to define them?
How would we deal with legacy code?

Observations

- ▶ Specifying objects intended behaviour as state machines is natural, but may be demanding and not easy to get right
- ▶ Stating, for each method, the required and ensured state of fields and parameters may be easier
Assertions are part of Java since 2006

Question

Can we get behavioural types from code with assertions?

The real challenge: infer usages

- ▶ Specify object protocols when writing OO code is not common practice.
- ▶ Is it reasonable to ask programmers to define them?
How would we deal with legacy code?

Observations

- ▶ Specifying objects intended behaviour as state machines is natural, but may be demanding and not easy to get right
- ▶ Stating, for each method, the required and ensured state of fields and parameters may be easier
Assertions are part of Java since 2006

Question

Can we get behavioural types from code with assertions?

The real challenge: infer usages

- ▶ Specify object protocols when writing OO code is not common practice.
- ▶ Is it reasonable to ask programmers to define them?
How would we deal with legacy code?

Observations

- ▶ Specifying objects intended behaviour as state machines is natural, but may be demanding and not easy to get right
- ▶ Stating, for each method, the required and ensured state of fields and parameters may be easier
Assertions are part of Java since 2006

Question

Can we get behavioural types from code with assertions?

The envisaged contribution

Infer, from OO code with assertions, behavioural (class) types ensuring safe interoperability

A type inference system: given a program

- ▶ either fails: the code is *not well-typed* (in the standard sense) or it may produce a run-time error due to calling methods in an *incorrect order*;
- ▶ or returns a new version of the code with the classes annotated with behavioural types, ensuring *object interoperability*.

Thanks!

The envisaged contribution

Infer, from OO code with assertions, behavioural (class) types ensuring safe interoperability

A type inference system: given a program

- ▶ either fails: the code is *not well-typed* (in the standard sense) or it may produce a run-time error due to calling methods in an *incorrect order*;
- ▶ or returns a new version of the code with the classes annotated with behavioural types, ensuring *object interoperability*.

Thanks!

The envisaged contribution

Infer, from OO code with assertions, behavioural (class) types ensuring safe interoperability

A type inference system: given a program

- ▶ either fails: the code is *not well-typed* (in the standard sense) or it may produce a run-time error due to calling methods in an *incorrect order*;
- ▶ or returns a new version of the code with the classes annotated with behavioural types, ensuring *object interoperability*.

Thanks!

The envisaged contribution

Infer, from OO code with assertions, behavioural (class) types ensuring safe interoperability

A type inference system: given a program

- ▶ either fails: the code is *not well-typed* (in the standard sense) or it may produce a run-time error due to calling methods in an *incorrect order*;
- ▶ or returns a new version of the code with the classes annotated with behavioural types, ensuring *object interoperability*.

Thanks!

The envisaged contribution

Infer, from OO code with assertions, behavioural (class) types ensuring safe interoperability

A type inference system: given a program

- ▶ either fails: the code is *not well-typed* (in the standard sense) or it may produce a run-time error due to calling methods in an *incorrect order*;
- ▶ or returns a new version of the code with the classes annotated with behavioural types, ensuring *object interoperability*.

Thanks!