

Message-Observing Sessions



Ryan Kavanagh and Brigitte Pientka

OOPSLA 2024



Université du Québec à Montréal and McGill University

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC).

Communicating systems are ubiquitous



Communicating systems are ubiquitous



Miscommunication is costly

- Incorrectly implemented communication protocols are costly

Miscommunication is costly

- Incorrectly implemented communication protocols are costly



Miscommunication is costly

- Incorrectly implemented communication protocols are costly



Miscommunication is costly

- Incorrectly implemented communication protocols are costly



Miscommunication is costly

- Incorrectly implemented communication protocols are costly



aws

yahoo!



Miscommunication is costly

- Incorrectly implemented communication protocols are costly



Canada Revenue
Agency

Agence du revenu
du Canada

Miscommunication is costly

- Incorrectly implemented communication protocols are costly



Canada Revenue
Agency

Agence du revenu
du Canada

- Want to **statically guarantee** our programs communicate correctly

Miscommunication is costly

- Incorrectly implemented communication protocols are costly



Canada Revenue
Agency

Agence du revenu
du Canada

- Want to statically guarantee our programs communicate correctly
- Want to **precisely specify** the desired communication behaviours

Key technology: Session types

Session types encode communication protocols as types.

Key technology: Session types

Session types encode communication protocols as types.

- **Multiparty session types** (MPST) specify systems top-down
 - + Specify rich interactions involving multiple processes
 - Typically not compositional, a closed-world approach

Key technology: Session types

Session types encode communication protocols as types.

- Multiparty session types (MPST) specify systems top-down
 - + Specify rich interactions involving multiple processes
 - Typically not compositional, a closed-world approach
- **Binary session types** specify systems bottom-up
 - Only specify local interactions between pairs of processes
 - + Compositional, with many expressive varieties

Key technology: Session types

Session types encode communication protocols as types.


- Multiparty session types (MPST) specify systems top-down
 - + Specify rich interactions involving multiple processes
 - Typically not compositional, a closed-world approach
- Binary session types specify systems bottom-up
 - Only specify local interactions between pairs of processes
 - + Compositional, with many expressive varieties

Research Problem

Can we have the best of both worlds:

the ability to specify **global properties** and **compositionality**?

Specifying communication-based concurrency



where

- P — process
- c_i — name of bidirectional communication channel
- A_i — protocol (session type) on channel c_i

Write $P [c_1 : A_1, \dots, c_n : A_n] (c_0 : A_0)$ to syntactically specify P

Bit streams

Bit stream protocol:

$$\text{bits} = (b_0 \text{ : bits}) \oplus (b_1 \text{ : bits})$$

Example communications satisfying `bits`:

$$\xrightarrow{\quad b_0 \ b_1 \ b_0 \ b_0 \ \dots \quad} c : \text{bits}$$

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 ⇒ send b1 on o; F(i; o)  
          | b1 ⇒ send b0 on o; F(i; o) }
```

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 ⇒ send b1 on o; F(i; o)  
          | b1 ⇒ send b0 on o; F(i; o) }
```

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 ⇒ send b1 on o; F(i; o)  
          | b1 ⇒ send b0 on o; F(i; o) }
```

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 ⇒ send b1 on o; F(i; o)  
          | b1 ⇒ send b0 on o; F(i; o) }
```

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 => send b1 on o; F(i; o)  
          | b1 => send b0 on o; F(i; o) }
```


Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 => send b1 on o; F(i; o)  
          | b1 => send b0 on o; F(i; o) }
```

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 => send b1 on o; F(i; o)  
          | b1 => send b0 on o; F(i; o) }
```

Specifying bit flipping



```
F [i : bits] (o : bits) =  
  case i { b0 ⇒ send b1 on o; F(i; o)  
          | b1 ⇒ send b0 on o; F(i; o) }
```

Problem

The specification `F [i : bits] (o : bits)` does not specify or enforce bit flipping!

Today's contribution: **Most**

Most is a language for compositionally specifying communication protocols that depend on communications occurring on other channels.

Today's contribution: **Most**

Most is a language for **compositionally** specifying communication protocols that depend on communications occurring on other channels.

Today's contribution: **Most**

Most is a language for compositionally specifying communication protocols that depend on communications occurring on other channels.

Today's contribution: **Most**

Most is a language for compositionally specifying communication protocols that depend on communications occurring on other channels.

Three ingredients:

1. **syntax** for specifying protocols with type-level computation
2. **semantic framework** for explaining protocols as sequences of allowed communications while taking dependency into account
3. **static typechecking** to ensure that processes satisfy their protocols

Most's Syntax

Introducing message-observing session types

$A, B := \dots$

$| (l \circlearrowleft A) \oplus (r \circlearrowleft B)$

other session types

labelled choice

Introducing message-observing session types

$A, B := \dots$

| $(l \circlearrowleft A) \oplus (r \circlearrowleft B)$

| $\text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$

other session types

labelled choice

label observation

Introducing message-observing session types

$A, B ::= \dots$

other session types

$| (l \circ A) \oplus (r \circ B)$

labelled choice

$| \text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$

label observation

Operational Intuition

$\text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$ reduces to A if l observed on channel c

$\text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$ reduces to B if r observed on channel c

Introducing message-observing session types

$A, B := \dots$

$| (l \circlearrowleft A) \oplus (r \circlearrowleft B)$

$| \text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$

other session types

labelled choice

label observation

Operational Intuition

$\text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$ reduces to A if l observed on channel c

$\text{CASE } c \{l \Rightarrow A \mid r \Rightarrow B\}$ reduces to B if r observed on channel c

See paper for how to observe termination and channel transmission!

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0 } 8 \text{ bits}) \oplus (\text{b1 } 8 \text{ bits})$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ \textasciitilde{ } bits}) \oplus (\text{b1} \text{ \textasciitilde{ } bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\begin{aligned} \text{bitsFlip}(i) = \text{CASE } i \{ & \text{b0} \Rightarrow (\text{b1} \text{ \textasciitilde{ } bitsFlip}(i)) \\ & | \text{b1} \Rightarrow (\text{b0} \text{ \textasciitilde{ } bitsFlip}(i)) \} \end{aligned}$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ \textasciitilde{ } bits}) \oplus (\text{b1} \text{ \textasciitilde{ } bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\text{bitsFlip}(i) = \text{CASE } i \left\{ \begin{array}{l} \text{b0} \Rightarrow (\text{b1} \text{ \textasciitilde{ } bitsFlip}(i)) \\ | \text{b1} \Rightarrow (\text{b0} \text{ \textasciitilde{ } bitsFlip}(i)) \end{array} \right\}$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ } \text{; bits}) \oplus (\text{b1} \text{ } \text{; bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\begin{aligned} \text{bitsFlip}(i) = \text{CASE } i \{ & \text{b0} \Rightarrow (\text{b1} \text{ } \text{; bitsFlip}(i)) \\ & | \text{b1} \Rightarrow (\text{b0} \text{ } \text{; bitsFlip}(i)) \} \end{aligned}$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ } \text{; bits}) \oplus (\text{b1} \text{ } \text{; bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\begin{aligned} \text{bitsFlip}(i) = \text{CASE } i \{ & \text{b0} \Rightarrow (\text{b1} \text{ } \text{; bitsFlip}(i)) \\ & | \text{b1} \Rightarrow (\text{b0} \text{ } \text{; bitsFlip}(i)) \} \end{aligned}$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ } \text{; bits}) \oplus (\text{b1} \text{ } \text{; bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\begin{aligned} \text{bitsFlip}(i) = \text{CASE } i \{ & \text{b0} \Rightarrow (\text{b1} \text{ } \text{; bitsFlip}(i)) \\ & | \text{b1} \Rightarrow (\text{b0} \text{ } \text{; bitsFlip}(i)) \} \end{aligned}$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ \# bits}) \oplus (\text{b1} \text{ \# bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\begin{aligned} \text{bitsFlip}(i) = \text{CASE } i \{ & \text{b0} \Rightarrow (\text{b1} \text{ \# bitsFlip}(i)) \\ & | \text{b1} \Rightarrow (\text{b0} \text{ \# bitsFlip}(i)) \} \end{aligned}$$

$$F [i : \text{bits}] (\text{o} : \text{bitsFlip}(i)) = \dots$$

Revisiting bit flipping

Bit stream protocol:

$$\text{bits} = (\text{b0} \text{ \# bits}) \oplus (\text{b1} \text{ \# bits})$$

Bit flipping protocol relative to a channel i : bits:

$$\begin{aligned} \text{bitsFlip}(i) = \text{CASE } i \{ & \text{b0} \Rightarrow (\text{b1} \text{ \# bitsFlip}(i)) \\ & | \text{b1} \Rightarrow (\text{b0} \text{ \# bitsFlip}(i)) \} \end{aligned}$$

$F [i : \text{bits}] (\text{o} : \text{bitsFlip}(i)) = \dots$

This technique can specify any stream transducer or multiplexer!


Insight #1: Type-level concurrent computation

First Key Insight

We can specify dependent communication protocols with a restricted form of **type-level concurrent computation**.


Composing processes

We want to be able to specify process compositions:



Composing processes


We want to be able to specify process compositions:



We want to reason about the whole in terms of the specifications of the parts.


Insight #2: Tracking ambient channels to achieve compositionality

The specification of the right process depends on an ambient i : bits:



Insight #2: Tracking ambient channels to achieve compositionality

The specification of the right process depends on an ambient i : bits:



We extend our specifications to track assumptions about ambient channels:

```
F {i : bits} [m : bitsFlip(i)] (o : idBits(i)) = ...
```

Rely-guarantee perspective ensures compositionality!

Most's Semantics

Denotational semantics

A **process** denotes a set of traces of messages sent or received on channels. For example, flipping bits received on m onto a channel o :

$$\llbracket F(m; o) \rrbracket = \{ \text{recv } b_0 \text{ on } m :: \text{send } b_1 \text{ on } o :: \dots, \\ \text{recv } b_1 \text{ on } m :: \text{send } b_0 \text{ on } o :: \dots, \dots \}$$

Denotational semantics

A **process** denotes a set of traces of messages sent or received on channels. For example, flipping bits received on m onto a channel o :

$$\llbracket F(m; o) \rrbracket = \{ \text{recv } b0 \text{ on } m :: \text{send } b1 \text{ on } o :: \dots, \\ \text{recv } b1 \text{ on } m :: \text{send } b0 \text{ on } o :: \dots, \dots \}$$

A **specification** denotes a set of allowed traces, interleaved with constraints on ambient channels. For example,

$$\llbracket \{i : \text{bits}\} [m : \text{bitsFlip}(i)] (o : \text{idBits}(i)) \rrbracket \\ = \{ \text{rely } b0 \text{ on } i :: \text{recv } b1 \text{ on } m :: \text{send } b0 \text{ on } o :: \dots, \dots \}$$

Typechecking Most

Constraint generation and checking

Typechecking P against a specification **generates constraints** \mathcal{T} on the what communications may appear on ambient channels:

$$P \Vdash \{AmbientCtx\} [ClientCtx] (a : A) // \mathcal{T}$$

Constraint generation and checking

Typechecking P against a specification generates constraints \mathcal{T} on the what communications may appear on ambient channels:

$$P \Vdash \{AmbientCtx\} [ClientCtx] (a : A) // \mathcal{T}$$

When typechecking process compositions, we check that each process satisfies the constraints that the other imposes on its channels.

Theorem

Our typechecking algorithm is semantically sound:

If P typechecks against a specification, then its traces are among those allowed by that specification.

Not the first dependent session type system, but crucially different:

- *Value-dependent session types*: types depend on transmitted values.
- *Label-dependent session types*: types depend on transmitted labels.

Multi-party session types provide a rich notion of process specification, but are quite complex and not compositional.

Future Work

- Introduce **sharing** to specify and verify shared services (databases, shared data structures, etc.)
- Develop an elegant **subtyping** relation to allow composition along channels with different types
- **Mechanize** the system and extract a verified compiler
- **Characterize the expressiveness gap** between Most and MPST

Takeaways

Concurrent type-level computation lets us compositionally specify protocols that vary based on ambient communications.

Most provides a significant step towards capturing message-dependent protocols and providing more precise specifications.

I am recruiting students! If this work sounds interesting, please come talk to me!