

Communication-Based Semantics for Recursive Session-Typed Processes

Ryan Kavanagh

September 28, 2021

Committee: Stephen Brookes, co-chair

Frank Pfenning, co-chair

Jan Hoffmann

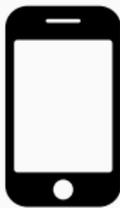
Luís Caires, Universidade Nova de Lisboa

Gordon Plotkin, University of Edinburgh

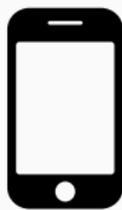
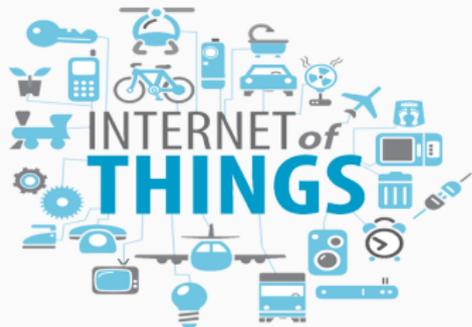
Communicating Systems and Session Types



Communicating Systems and Session Types



Communicating Systems and Session Types



Programs written in **session-typed programming languages** are **guaranteed** to obey their protocols.

Key Technique: Program Equivalence

“Program equivalence is arguably one of the most interesting and at the same time important problems in formal verification.”¹

¹Lahiri et. al. “Program Equivalence”, Dagstuhl Reports, Vol. 8, Issue 4, pp. 1–19, 2018.

Some Existing Approaches to Equivalence

There are existing notions of program equivalence for session-typed languages:

Some Existing Approaches to Equivalence

There are existing notions of program equivalence for session-typed languages:

- Wadler's Classical Processes (CP): Atkey [2017] gives a relational semantics.
- Hypersequent CP: Kokke et al. [2019] give a denotational semantics using Brzozowski derivatives.

Some Existing Approaches to Equivalence

There are existing notions of program equivalence for session-typed languages:

- Wadler's Classical Processes (CP): Atkey [2017] gives a relational semantics.
- Hypersequent CP: Kokke et al. [2019] give a denotational semantics using Brzozowski derivatives.
- Synchronous session-typed π -calculus: Castellan and Yoshida [2019] give a game semantics.

Some Existing Approaches to Equivalence

There are existing notions of program equivalence for session-typed languages:

- Wadler's Classical Processes (CP): Atkey [2017] gives a relational semantics.
- Hypersequent CP: Kokke et al. [2019] give a denotational semantics using Brzozowski derivatives.
- Synchronous session-typed π -calculus: Castellan and Yoshida [2019] give a game semantics.

Some Existing Approaches to Equivalence

There are existing notions of program equivalence for session-typed languages:

- Wadler's Classical Processes (CP): Atkey [2017] gives a relational semantics.
- Hypersequent CP: Kokke et al. [2019] give a denotational semantics using Brzozowski derivatives.
- Synchronous session-typed π -calculus: Castellan and Yoshida [2019] give a game semantics.

Problem: It is not clear how to extend these approaches to handle full-featured languages.

Reasoning About Programs in Rich Languages

When one attempts to combine language concepts, unexpected and counterintuitive interactions arise. At this point, even the most experienced designer's intuition must be buttressed by a rigorous definition of what the language means. — John Reynolds, 1990

Reasoning About Programs in Rich Languages

We want to reason about programs in a session-typed language with:

- general recursion at the program and type level
- functional programming features
- higher-order features: send/receive channels and programs

Back to Fundamentals

A **process** is a computational agent that interacts with its environment solely through communication.

Communication is a sequence of atomic observable events caused by a process.

The Key Premise

Communication is the only observable phenomenon of processes!

Thesis Statement

Communication-based semantics elucidate the structure of session-typed languages and allow us to reason about programs written in these languages.

Our Laboratory: Polarized SILL

We will study “Polarized SILL”, a language with:

1. a functional programming layer
2. session-typed message passing concurrency
3. general recursion (types and programs)
4. higher-order features: processes can send/receive channels and programs

Contributions

We give Polarized SILL

1. **An observed communication semantics**
2. **A communication-based testing equivalences framework**
3. **A communication-based denotational semantics**

and we use these semantics to reason about processes.

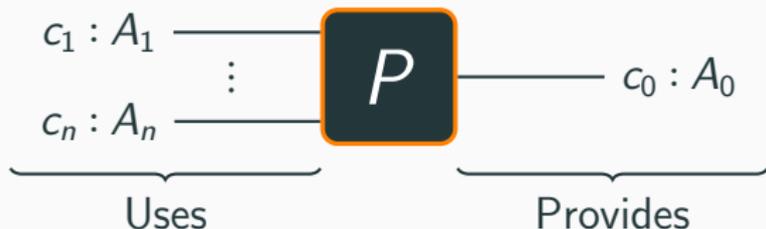
Polarized SILL



Where

- c_i — channel name
- A_i — protocol (session type) for channel c_i
- P — process

Polarized SILL



Abbreviate as:

$$c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0 \quad (n \geq 0)$$
$$\Delta \vdash P :: c_0 : A_0$$

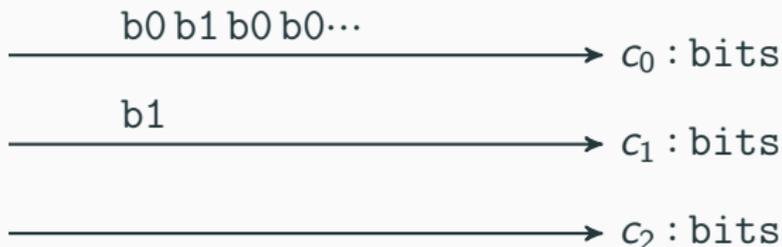
where $\Delta = c_1 : A_1, \dots, c_n : A_n$.

Bit Streams in SILL

Bit stream protocol:

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

Example communications satisfying bits:

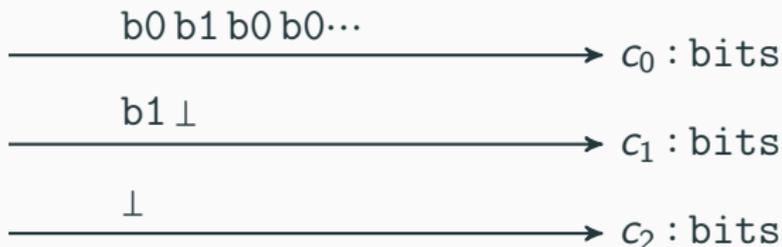


Bit Streams in SILL

Bit stream protocol:

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

Example communications satisfying bits:



Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                    | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                    | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                    | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                    | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Flipping Bits



```
i : bits |- F :: o : bits
```

```
o <- F <- i = case i { b0 => o.b1; o <- F <- i  
                      | b1 => o.b0; o <- F <- i }
```

Observed Communication Semantics

Idea: The meaning of a process is the communications we observe during its execution.

Idea: The meaning of a process is the communications we observe during its execution.

Questions:

1. What are observed communications?

Idea: The meaning of a process is the communications we observe during its execution.

Questions:

1. What are observed communications?
2. How do we observe them?

Session-Typed Communications

A session type specifies permitted communications.

Session-Typed Communications

A session type specifies permitted communications.

Write $w \varepsilon A$ to mean w is a communication satisfying the session type A .

Session-Typed Communications

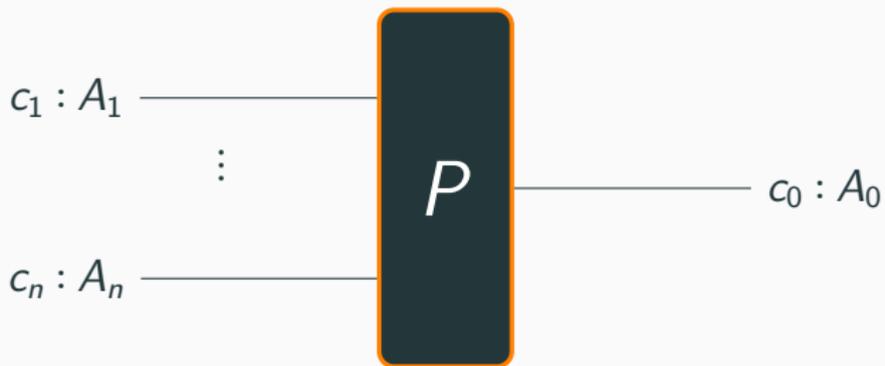
A session type specifies permitted communications.

Write $w \varepsilon A$ to mean w is a communication satisfying the session type A .

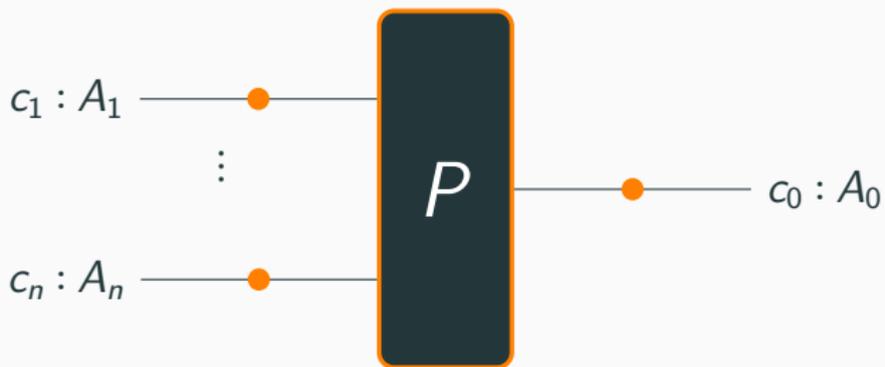
Examples:

- The **empty communication** $\perp \varepsilon A$.
- **Bit stream communications** are $(b0, w) \varepsilon \text{bits}$ and $(b1, w) \varepsilon \text{bits}$ where $w \varepsilon \text{bits}$.

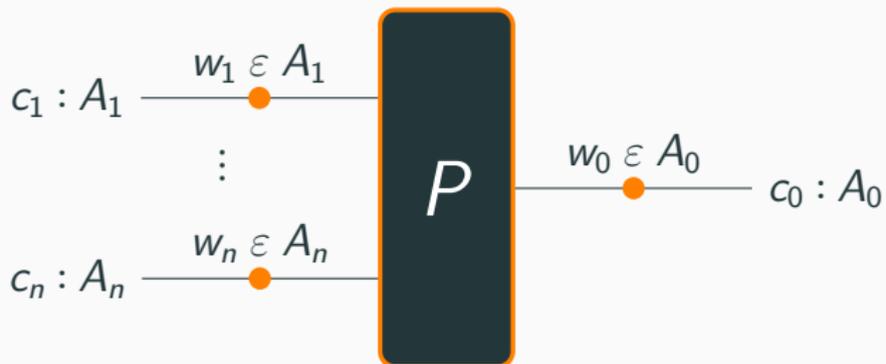
Observing Communications



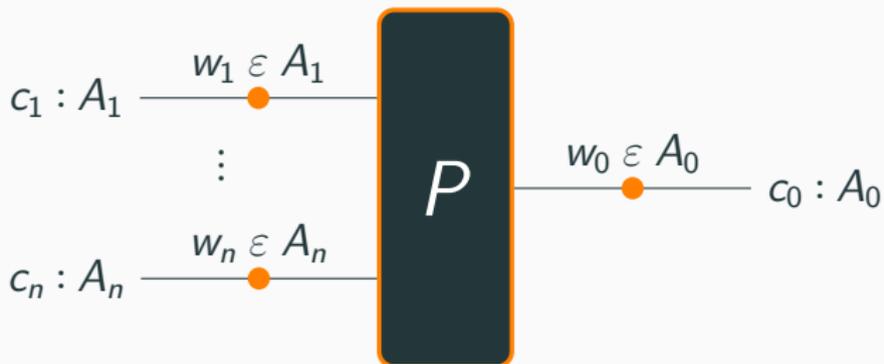
Observing Communications



Observing Communications



Observing Communications



$$\langle c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0 \rangle_{c_0, \dots, c_n} = (c_0 : w_0, \dots, c_n : w_n).$$

Example Observed Communications

Consider the process S sending a stream of zero bits:

```
⊢ S :: i : bits
```

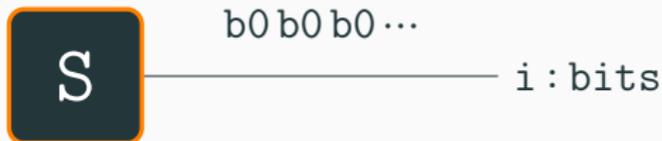
```
i <- S = i.b0; i <- S
```

Example Observed Communications

Consider the process S sending a stream of zero bits:

```
⊢ S :: i : bits
```

```
i <- S = i.b0; i <- S
```

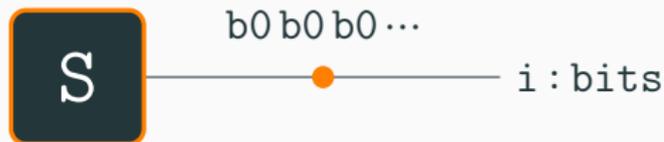


Example Observed Communications

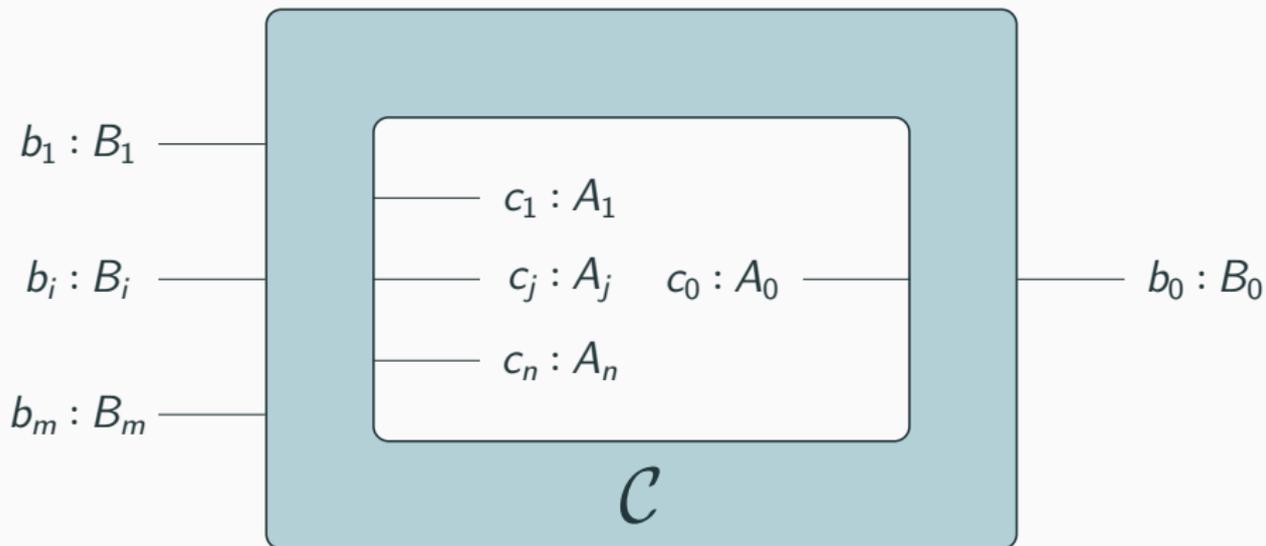
Consider the process S sending a stream of zero bits:

```
⊢ S :: i : bits
```

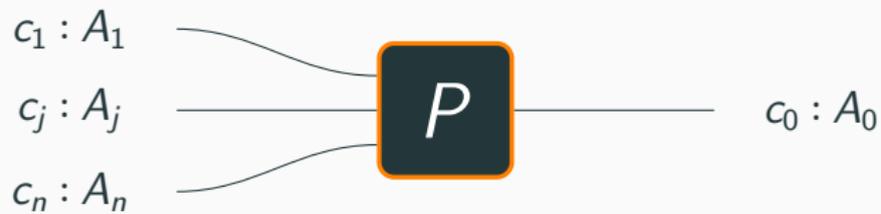
```
i <- S = i.b0; i <- S
```


$$\langle \vdash S :: i : bits \rangle_i = (i : (b0, (b0, (b0, \dots))))$$

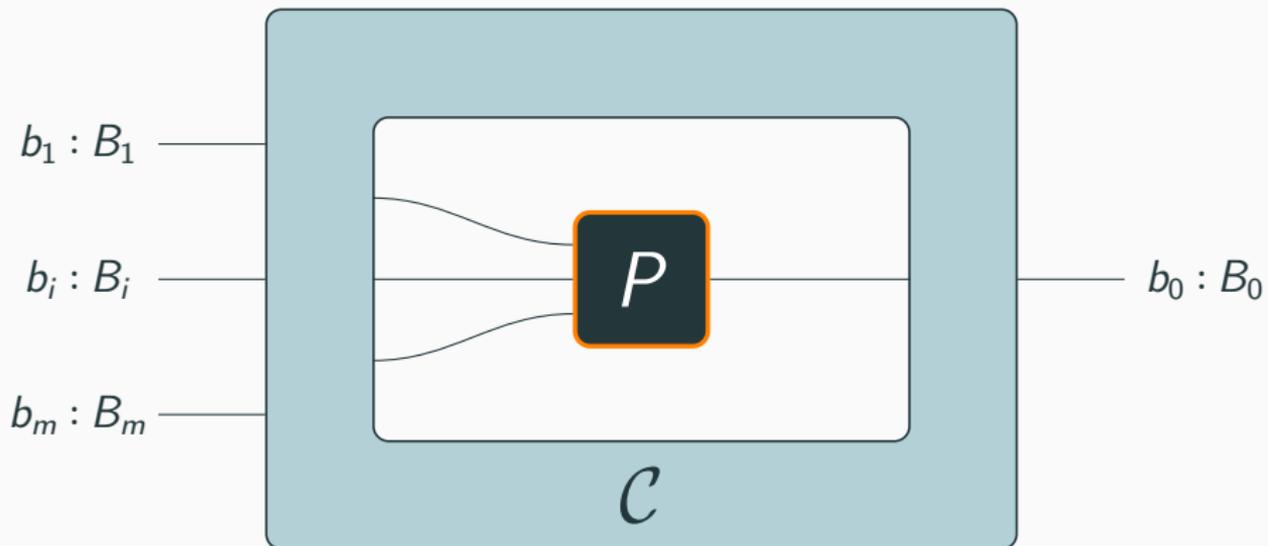
Observing Communication Between Processes



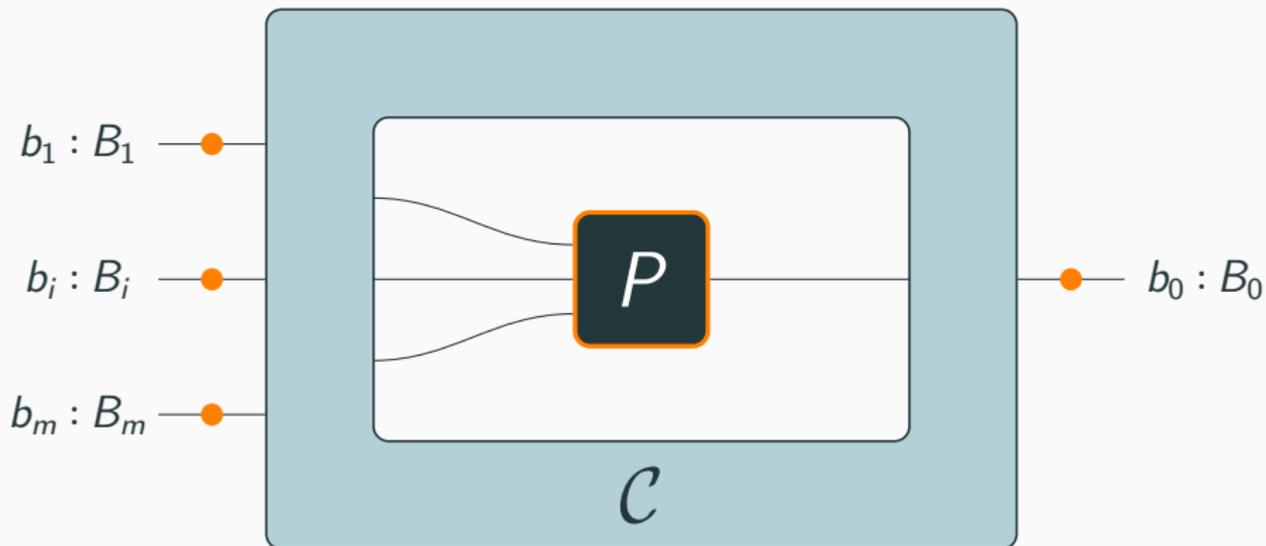
Observing Communication Between Processes



Observing Communication Between Processes

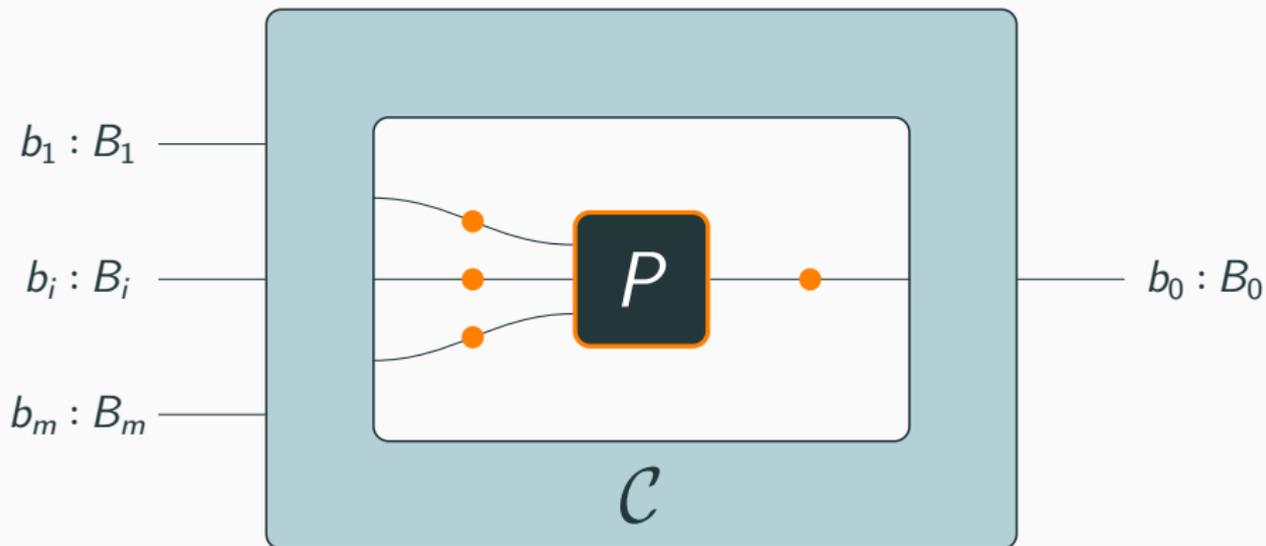


Observing Communication Between Processes



$$\langle b_1 : B_1, \dots, b_m : B_m \vdash C[P] :: b_0 : B_0 \rangle_{b_0, \dots, b_m} = (b_0 : w_0, \dots, b_m : w_m)$$

Observing Communication Between Processes



$$\langle b_1 : B_1, \dots, b_m : B_m \vdash C[P] :: b_0 : B_0 \rangle_{b_0, \dots, b_m} = (b_0 : w_0, \dots, b_m : w_m)$$

$$\langle b_1 : B_1, \dots, b_m : B_m \vdash C[P] :: b_0 : B_0 \rangle_{c_0, \dots, c_n} = (c_0 : w'_0, \dots, c_n : w'_n)$$

More Example Observed Communications



More Example Observed Communications



More Example Observed Communications



More Example Observed Communications


$$\langle \vdash \mathcal{C}[F] :: o : \text{bits} \rangle_o = (o : (b1, (b1, (b1, \dots))))$$

More Example Observed Communications


$$\langle \vdash \mathcal{C}[F] :: o : \text{bits} \rangle_o = (o : (b1, (b1, (b1, \dots))))$$
$$\langle \vdash \mathcal{C}[F] :: o : \text{bits} \rangle_i = (i : (b0, (b0, (b0, \dots))))$$

More Example Observed Communications


$$\langle \vdash C[F] :: o : \text{bits} \rangle_o = (o : (b1, (b1, (b1, \dots))))$$
$$\langle \vdash C[F] :: o : \text{bits} \rangle_i = (i : (b0, (b0, (b0, \dots))))$$
$$\langle \vdash C[F] :: o : \text{bits} \rangle_{i,o} = (i : (b0, \dots), o : (b1, \dots))$$

Fairness



Fairness



Fairness



Fairness



$$\langle \vdash C[F] :: o : \text{bits} \rangle_{i,o} = (i : (b_0, (b_0, \dots)), o : \perp)$$

Fairness



$$\langle \vdash C[F] :: o : \text{bits} \rangle_{i,o} = (i : (b_0, (b_0, \dots)), o : \perp)$$

Theorem

Observed communications are independent of the choice of fair execution.

Contributions

We give Polarized SILL

1. **An observed communication semantics**
2. **A communication-based testing equivalences framework**
3. **A communication-based denotational semantics**

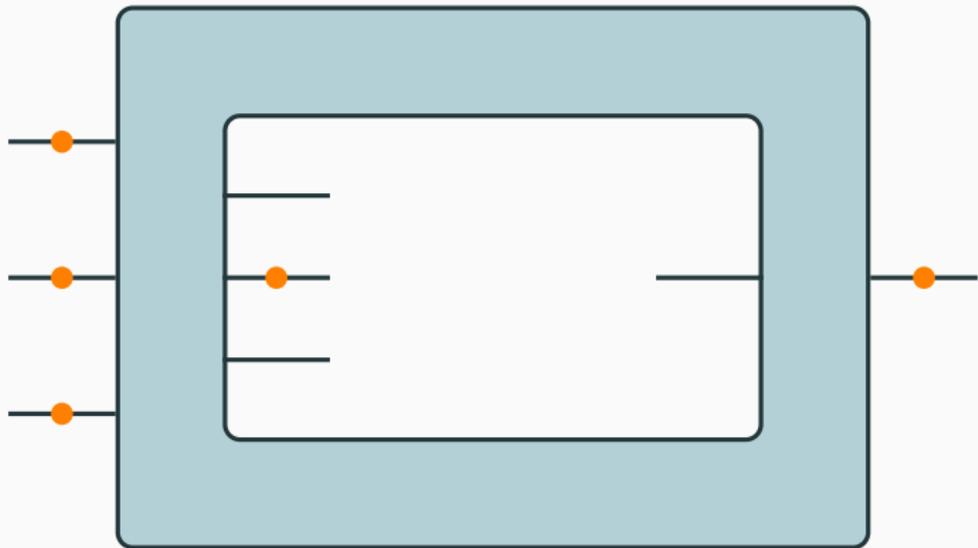
and we use these semantics to reason about processes.

Communication-Based Testing Equivalences

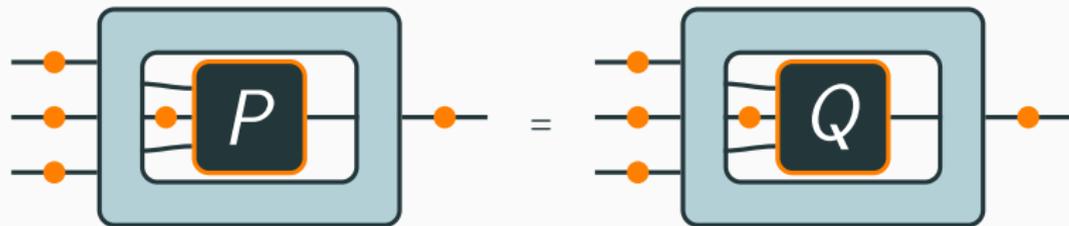
Testing Equivalences

Main Idea: Two processes are **equivalent** if we cannot observe any differences through experimentation.

Communication Experiments



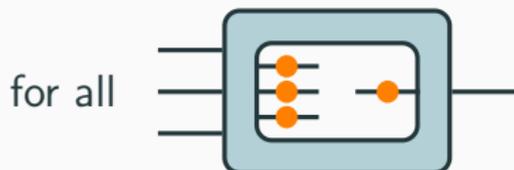
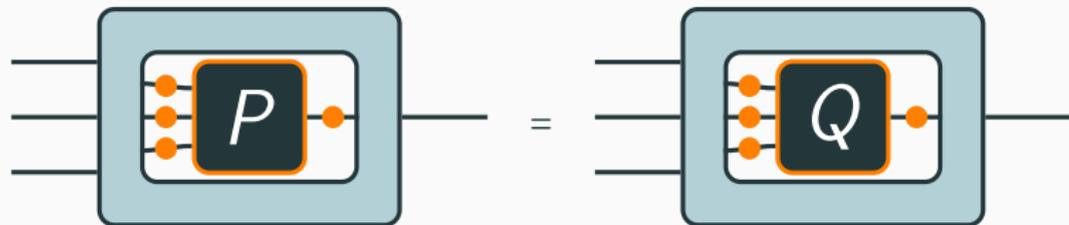
Performing Experiments



Internal Communication Equivalence

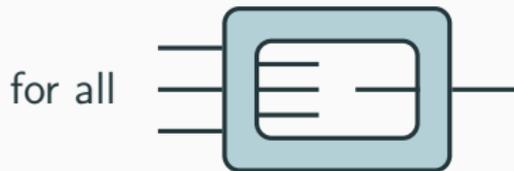


are **internally communication equivalent** if



Congruence Relations

An equivalence relation \equiv is a **congruence** if



Not A Congruence

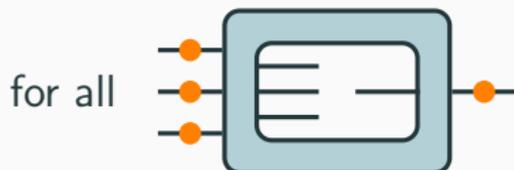
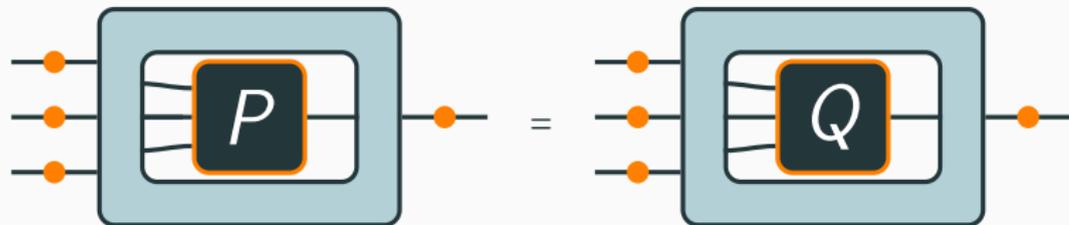
Theorem

Internal communication equivalence is not a congruence relation.

External Communication Equivalence



are **externally communication equivalent** if



Theorem

External communication equivalence is a congruence relation.

Theorem

External communication equivalence is a congruence relation.

Barbed congruence is the canonical notion of process equivalence.

Properties of External Communication Equivalence

Theorem

External communication equivalence is a congruence relation.

Barbed congruence is the canonical notion of process equivalence.

Theorem

Processes are external communication equivalent if and only if they are barbed congruent.

New Semantics, Same Refrain

“Processes are equivalent if [...] for all



Contributions

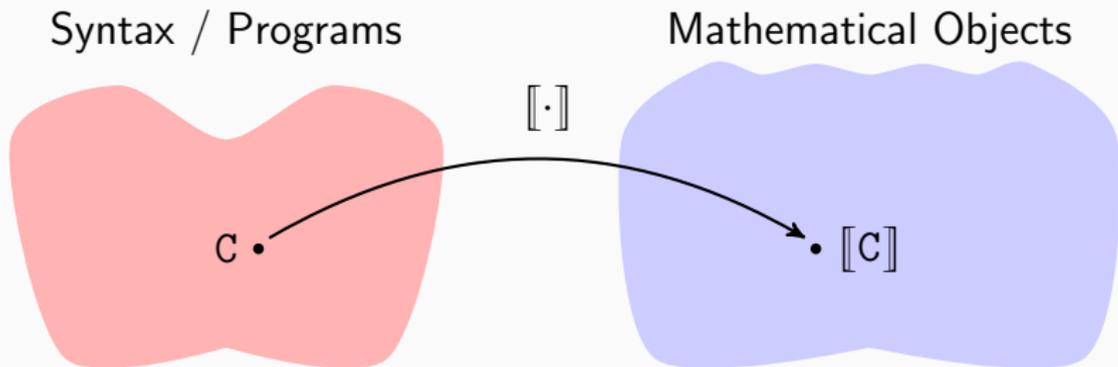
We give Polarized SILL

1. **An observed communication semantics**
2. **A communication-based testing equivalences framework**
3. **A communication-based denotational semantics**

and we use these semantics to reason about processes.

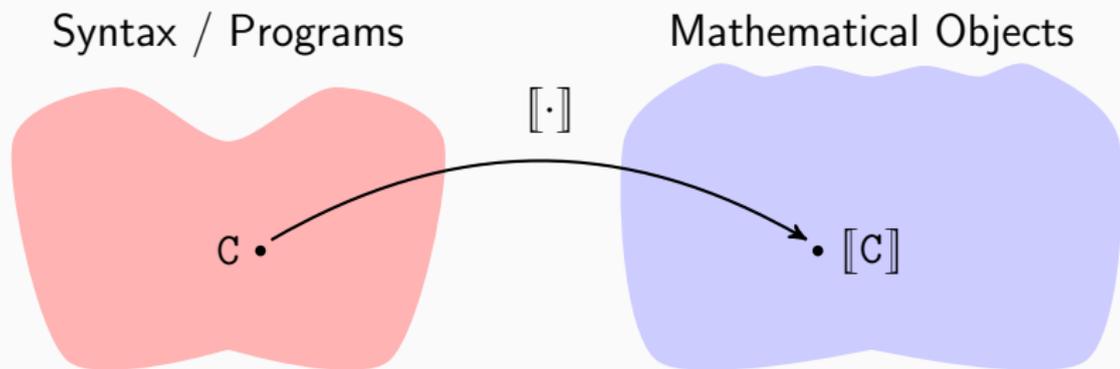
Denotational Semantics

The Denotational Approach



Compositional: the meaning of a program is a function of the meanings of its parts.

The Denotational Approach



Compositional: the meaning of a program is a function of the meanings of its parts.

Programs C and C' are **semantically equivalent** if $[[C]] = [[C']]$.

Denoting Protocols and Processes

A **protocol** A denotes a **domain** $\llbracket A \rrbracket$ of permissible communications.

A **process** $c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0$ denotes a **continuous function** $\llbracket P \rrbracket : \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \rightarrow \llbracket A_0 \rrbracket$.

Monotonicity

Significance: “New” input does not affect “old” output.

If



then never



Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.

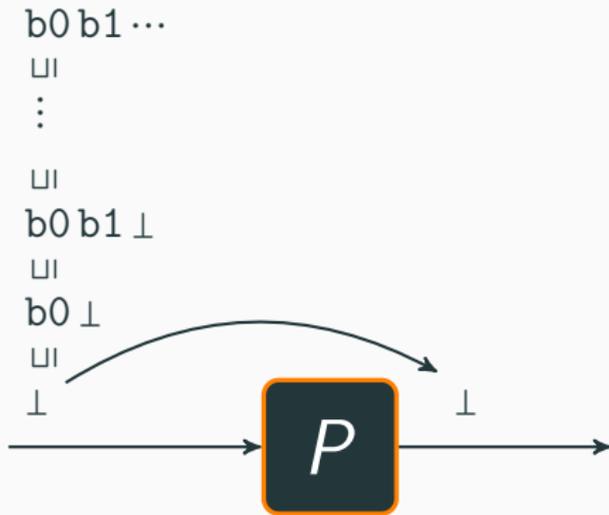
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.

```
b0 b1 ...  
⊔  
⋮  
⊔  
b0 b1 ⊥  
⊔  
b0 ⊥  
⊔  
⊥
```

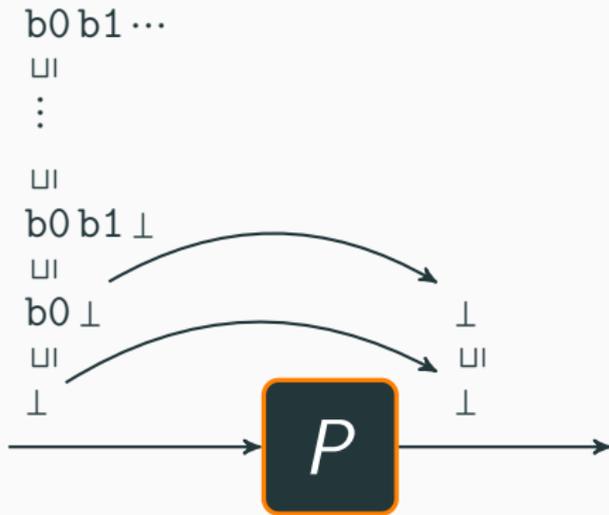
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.



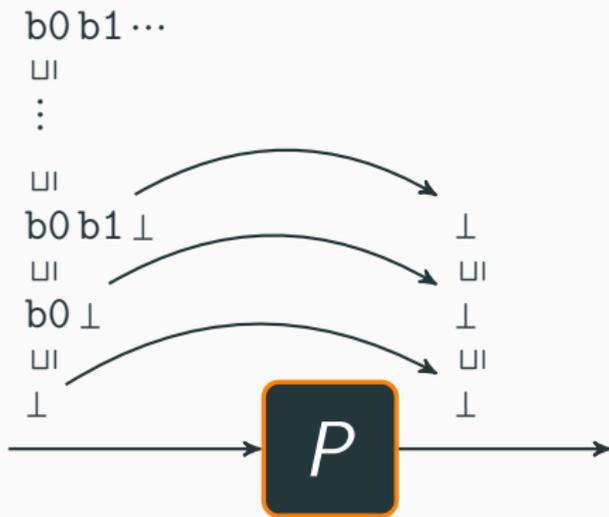
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.



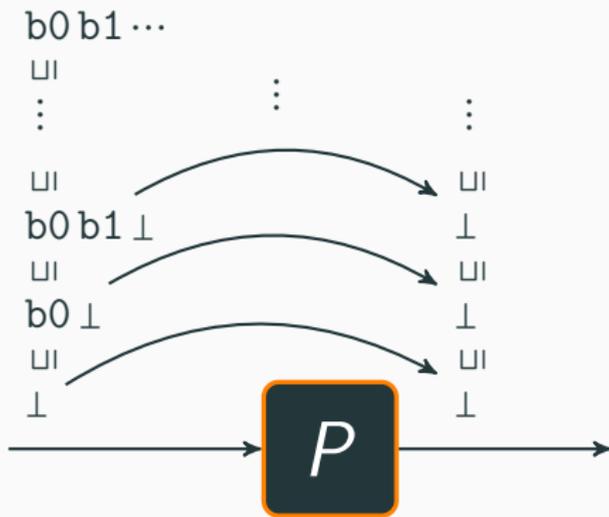
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.



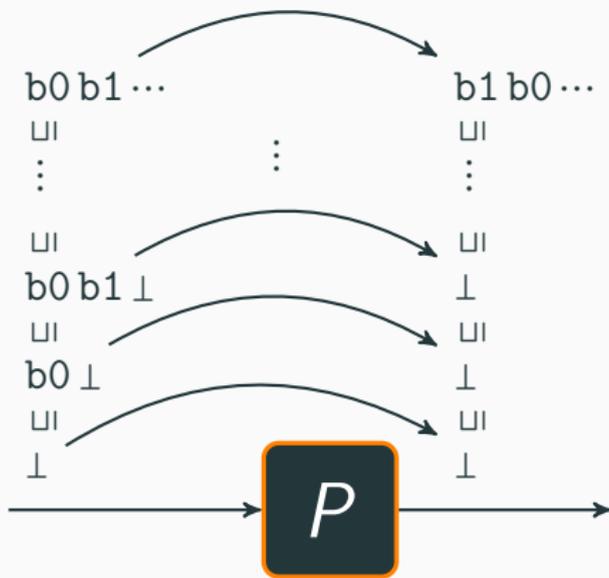
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.



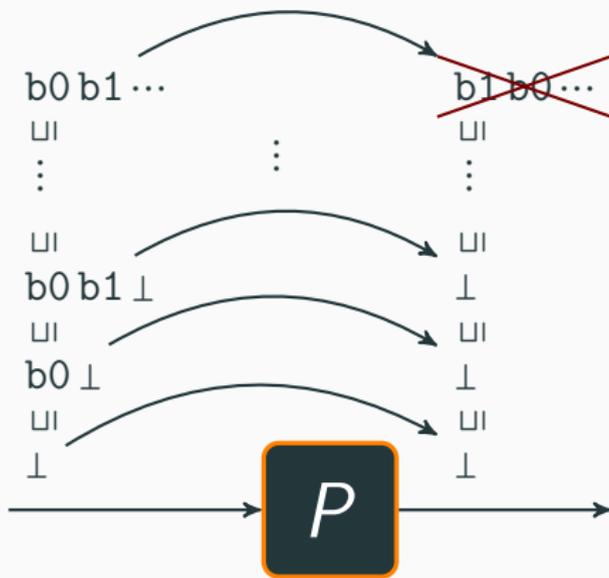
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.



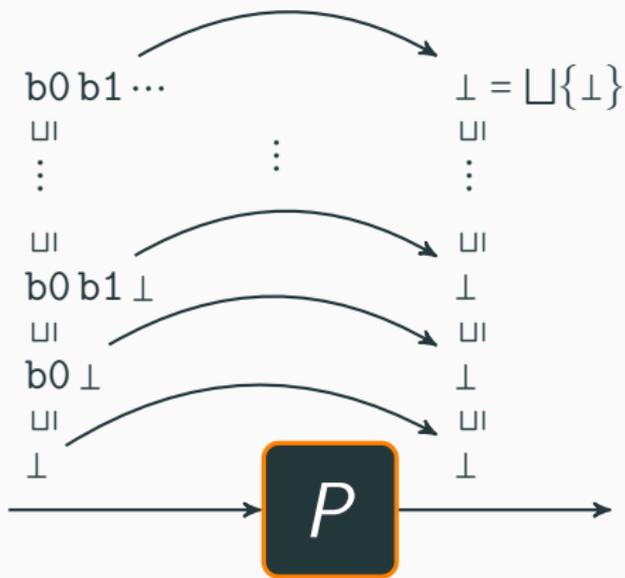
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.



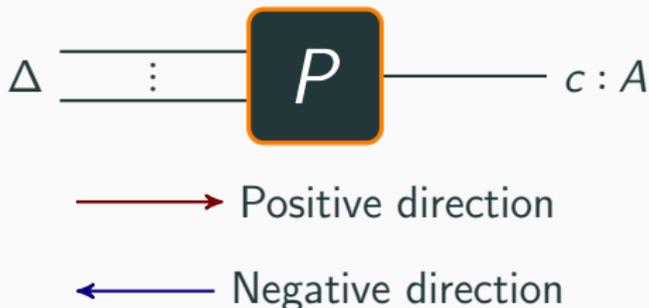
Continuity

Slogan: Processes cannot decide to send output only after observing entire infinite inputs.

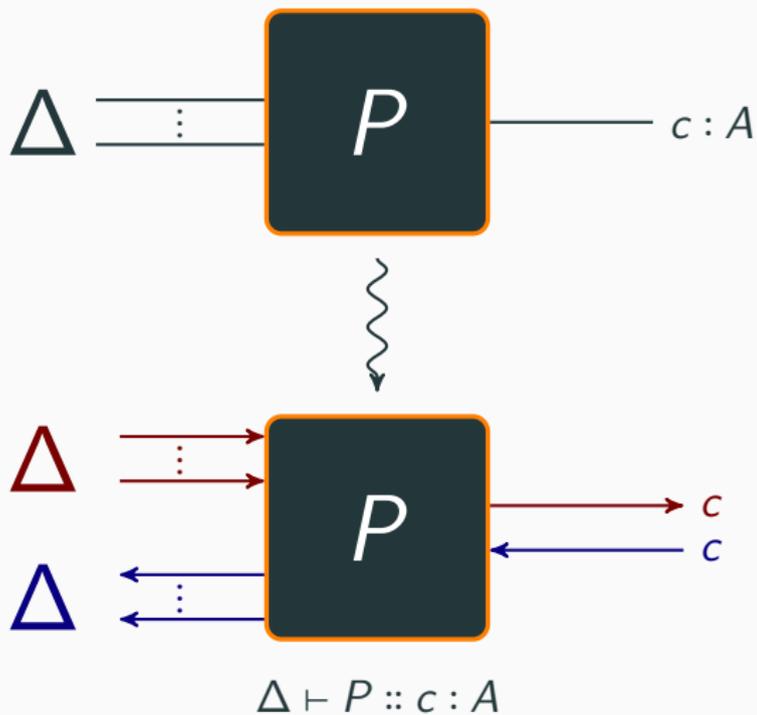


Polarity of Communication

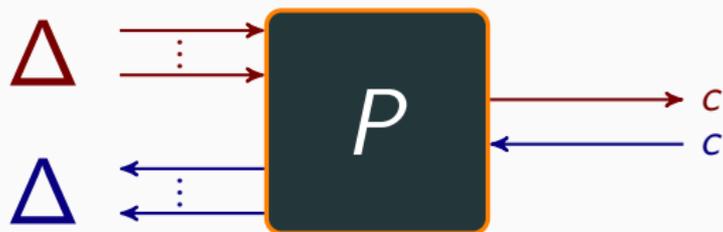
The **polarity** of a protocol is the direction in which its messages flow on channels.



Splitting Channels



Splitting Channels



$\Delta \vdash P :: c : A$

$\llbracket P \rrbracket : \Delta \times c \rightarrow \Delta \times c$

The Bidirectional Version

A **protocol** A denotes the domains

- $\llbracket A \rrbracket$ of negative (right-to-left) communications, and
- $\llbracket A \rrbracket$ of positive (left-to-right) communications.

A **process** $c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0$ denotes a continuous function

$$\begin{aligned} \llbracket P \rrbracket : \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \times \llbracket A_0 \rrbracket &\rightarrow \\ &\rightarrow \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \times \llbracket A_0 \rrbracket \end{aligned}$$

Decomposing Communications

A **protocol** A denotes a decomposition function

$$\langle A \rangle : \llbracket A \rrbracket \rightarrow \llbracket A \rrbracket \times \llbracket A \rrbracket$$

from the domain $\llbracket A \rrbracket$ of complete communications into the domains

- $\llbracket A \rrbracket$ of positive (left-to-right) communications,
- $\llbracket A \rrbracket$ of negative (right-to-left) communications.

Processes “Fill In” Partial Communications

A **process** $c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0$ denotes a continuous function

$$\begin{aligned} \llbracket P \rrbracket : \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \times \llbracket A_0 \rrbracket &\rightarrow \\ &\rightarrow \llbracket A_1 \rrbracket \times \dots \llbracket A_n \rrbracket \times \llbracket A_0 \rrbracket \end{aligned}$$

that is compatible with the decompositions

$$\langle A_i \rangle : \llbracket A_i \rrbracket \rightarrow \llbracket A_i \rrbracket \times \llbracket A_i \rrbracket.$$

The Functional Layer

- Simply-typed λ -calculus with a fixed-point operator
- Typing judgment: $\Psi \Vdash M : \tau$

The Functional Layer

- Simply-typed λ -calculus with a fixed-point operator
- Typing judgment: $\Psi \Vdash M : \tau$
- Standard denotational semantics:

$$\llbracket x_1 : \tau_1, \dots, x_n : \tau_n \Vdash M : \tau \rrbracket : \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \rightarrow \llbracket \tau \rrbracket$$

The Functional Layer

- Simply-typed λ -calculus with a fixed-point operator
- Typing judgment: $\Psi \Vdash M : \tau$
- Standard denotational semantics:

$$\llbracket x_1 : \tau_1, \dots, x_n : \tau_n \Vdash M : \tau \rrbracket : \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \rightarrow \llbracket \tau \rrbracket$$

The Functional Layer

- Simply-typed λ -calculus with a fixed-point operator
- Typing judgment: $\Psi \Vdash M : \tau$
- Standard denotational semantics:

$$\llbracket x_1 : \tau_1, \dots, x_n : \tau_n \Vdash M : \tau \rrbracket : \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \rightarrow \llbracket \tau \rrbracket$$

- Includes quoted processes as a base type

The Functional Layer

- Simply-typed λ -calculus with a fixed-point operator
- Typing judgment: $\Psi \Vdash M : \tau$
- Standard denotational semantics:

$$\llbracket x_1 : \tau_1, \dots, x_n : \tau_n \Vdash M : \tau \rrbracket : \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \rightarrow \llbracket \tau \rrbracket$$

- Includes quoted processes as a base type

Processes can depend on functional values through contexts Ψ :

$$\Psi; c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0$$

The Functional Layer

- Simply-typed λ -calculus with a fixed-point operator
- Typing judgment: $\Psi \Vdash M : \tau$
- Standard denotational semantics:

$$\llbracket x_1 : \tau_1, \dots, x_n : \tau_n \Vdash M : \tau \rrbracket : \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \rightarrow \llbracket \tau \rrbracket$$

- Includes quoted processes as a base type

Processes can depend on functional values through contexts Ψ :

$$\Psi; c_1 : A_1, \dots, c_n : A_n \vdash P :: c_0 : A_0$$

Processes now denote continuous functions

$$\begin{aligned} \llbracket P \rrbracket : \llbracket \Psi \rrbracket \rightarrow \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \times \llbracket A_0 \rrbracket \rightarrow \\ \rightarrow \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \times \llbracket A_0 \rrbracket \end{aligned}$$

Recall that processes P and Q are **denotationally equivalent** if $\llbracket P \rrbracket = \llbracket Q \rrbracket$.

Recall that processes P and Q are **denotationally equivalent** if $\llbracket P \rrbracket = \llbracket Q \rrbracket$.

Theorem

If two processes are denotationally equivalent, then they are external communication equivalent and barbed congruent.

Contributions

We give Polarized SILL

1. **An observed communication semantics**
2. **A communication-based testing equivalences framework**
3. **A communication-based denotational semantics**

and we use these semantics to reason about processes.

Thesis Statement

Communication-based semantics elucidate the structure of session-typed languages and allow us to reason about programs written in these languages.

Other Results

1. Modelling recursive types required new techniques for reasoning about parametrized fixed points of functors [MFPS'20]

Other Results

1. Modelling recursive types required new techniques for reasoning about parametrized fixed points of functors [MFPS'20]
2. A study of fairness for multiset rewriting systems [EXPRESS/SOS'20]

Other Results

1. Modelling recursive types required new techniques for reasoning about parametrized fixed points of functors [MFPS'20]
2. A study of fairness for multiset rewriting systems [EXPRESS/SOS'20]
3. A collection of case studies to which I apply these techniques

Future Work

1. Applications to richer protocols, e.g., dependent protocols
2. Applications to richer communication topologies, e.g., multicast

Acknowledgements



Thesis Statement

Communication-based semantics elucidate the structure of session-typed languages and allow us to reason about programs written in these languages.

Backup Slides

Relation to Deterministic Networks

My semantics generalizes Kahn's 1974 semantics for deterministic networks to support:

1. session-typed communication instead of streams of values of simple type like integers or booleans
2. *bidirectional* communication instead of unidirectional streams of values

Generalizing Kahn-style semantics to handle non-determinism is difficult because of the Keller and Brock-Ackerman anomalies. Though execution in Polarized SILL is non-deterministic, its processes have deterministic input/output behaviour.

Relation to the Geometry of Interaction I

My semantics exists in a Gol construction $\mathcal{G}(\mathbf{CPO})$:

- Objects are pairs (A^+, A^-) of objects A^+, A^- from **CPO**
- Morphisms $f : (A^+, A^-) \rightarrow (B^+, B^-)$ are morphism $\hat{f} : A^+ \times B^- \rightarrow A^- \times B^+$ in **CPO**
- Composition $g \circ f$ is $\text{Tr}(\hat{g} \times \hat{f})$

Expressing my semantics in this construction:

$$\begin{aligned} & \llbracket \Delta_1, \Delta_2 \vdash c \leftarrow P; Q :: d : D \rrbracket \\ &= \llbracket \Delta_2, c : C \vdash Q :: d : D \rrbracket \circ \llbracket \Delta_1 \vdash P :: c : C \rrbracket \end{aligned}$$

Relation to the Geometry of Interaction II

- Abramsky and Jagadeesan (1994) use this construction to give a type-free interpretation of classical linear logic where all types denote the same “universal domain”
- Abramsky, Haghverdi, and Scott (2002) use it to give an algebraic framework for Girard’s Geometry of Interaction
- I use it to give a semantics that captures the computational aspects of a programming language with recursion

Relation to Atkey's Denotational Semantics

In Atkey's denotational semantics for CP:

- Protocols denote sets of communications
- $\llbracket \vdash P :: \Gamma \rrbracket \subset \llbracket \Gamma \rrbracket$ is a relation containing the possible observed communications on its free channels, e.g.,

$$\llbracket \vdash x \leftrightarrow y :: x : A, y : A^\perp \rrbracket = \{(a, a) \mid a \in \llbracket A \rrbracket\}$$

$$\llbracket \mathbf{1} \rrbracket = \llbracket \mathbf{1}^\perp \rrbracket = \{*\}$$

$$\llbracket \vdash x[] :: x : \mathbf{1} \rrbracket = \{(*)\}$$

$$\llbracket \vdash x().P :: \Gamma, x : \mathbf{1}^\perp \rrbracket = \{(\gamma, *) \mid \gamma \in \llbracket \vdash P :: \Gamma \rrbracket\}$$

$$\llbracket \vdash \nu x.(P|Q) :: \Gamma, \Delta \rrbracket = \{(\gamma, \delta) \mid (\gamma, a) \in \llbracket \vdash P :: \Gamma, x : A \rrbracket, \\ (\delta, a) \in \llbracket \vdash Q :: \Delta, x : A^\perp \rrbracket\}$$