

CS 4110

Programming Languages & Logics

Lecture 36
Concurrency

3 December 2014



Announcements

- None!

Concurrency

All of the languages we have seen so far in this course have been sequential, performing one step of computation at a time.

In the next few lectures we will consider languages where multiple threads of execution may be interleaved simultaneously.

These languages can be used to model computations that execute on parallel and distributed architectures.

IMP with Parallel Composition

As a first step, let's extend IMP with a new a parallel composition command:

IMP with Parallel Composition

As a first step, let's extend IMP with a new a parallel composition command:

$a ::= x \mid n \mid a_1 + a_2$

$b ::= \mathbf{true} \mid \mathbf{false} \mid a_1 < a_2$

$c ::= \mathbf{skip} \mid x := a \mid c_1; c_2 \mid \mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2 \mid \mathbf{while } b \mathbf{ do } c$
 $\mid c_1 \parallel c_2$

Operational Semantics

and extend the small-step operational semantics with the following rules for $c_1 \parallel c_2$, which interleave the execution of c_1 and c_2 :

$$\frac{\langle \sigma, c_1 \rangle \rightarrow \langle \sigma', c'_1 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c_2 \rangle}$$

Operational Semantics

and extend the small-step operational semantics with the following rules for $c_1 \parallel c_2$, which interleave the execution of c_1 and c_2 :

$$\frac{\langle \sigma, c_1 \rangle \rightarrow \langle \sigma', c'_1 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c_2 \rangle}$$

$$\frac{\langle \sigma, c_2 \rangle \rightarrow \langle \sigma', c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c_1 \parallel c'_2 \rangle}$$

Operational Semantics

and extend the small-step operational semantics with the following rules for $c_1 \parallel c_2$, which interleave the execution of c_1 and c_2 :

$$\frac{\langle \sigma, c_1 \rangle \rightarrow \langle \sigma', c'_1 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c_2 \rangle}$$

$$\frac{\langle \sigma, c_2 \rangle \rightarrow \langle \sigma', c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c_1 \parallel c'_2 \rangle}$$

$$\langle \sigma, \mathbf{skip} \parallel \mathbf{skip} \rangle \rightarrow \langle \sigma, \mathbf{skip} \rangle$$

Operational Semantics

and extend the small-step operational semantics with the following rules for $c_1 \parallel c_2$, which interleave the execution of c_1 and c_2 :

$$\frac{\langle \sigma, c_1 \rangle \rightarrow \langle \sigma', c'_1 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c_2 \rangle}$$

$$\frac{\langle \sigma, c_2 \rangle \rightarrow \langle \sigma', c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c_1 \parallel c'_2 \rangle}$$

$$\langle \sigma, \mathbf{skip} \parallel \mathbf{skip} \rangle \rightarrow \langle \sigma, \mathbf{skip} \rangle$$

Note that the rules for parallel compositions $c_1 \parallel c_2$ allow either sub-command to take a step; two sub-commands can interleave read and write operations involving the same store.

Process Calculi

In the 1970s, Tony Hoare, Robin Milner, and others correctly observed that in the future, computers would have multiple computing cores, but each would have its own independent store.

Hoare's Communicating Sequential Processes were an early and highly-influential language that capture a *message passing* form of concurrency.

Many languages have built on CSP including Milner's CCS and π -calculus, Petri nets, and others.

π -calculus Syntax

The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

π -calculus Syntax

The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

The key constructs are based on the ability to interact by sending and receiving channel names

π -calculus Syntax

The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

The key constructs are based on the ability to interact by sending and receiving channel names

$$x, y, z \in \mathcal{N}$$

Names

π -calculus Syntax

The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

The key constructs are based on the ability to interact by sending and receiving channel names

$$\begin{array}{ll} x, y, z \in \mathcal{N} & \text{Names} \\ \pi ::= \tau \mid \bar{x}(y) \mid x(y) \mid [x = y] \pi & \text{Prefixes} \end{array}$$

π -calculus Syntax

The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

The key constructs are based on the ability to interact by sending and receiving channel names

$x, y, z \in \mathcal{N}$	<i>Names</i>
$\pi ::= \tau \mid \bar{x}(y) \mid x(y) \mid [x = y] \pi$	<i>Prefixes</i>
$M, N ::= \mathbf{0} \mid \pi.P \mid M + M$	<i>Summations</i>

π -calculus Syntax

The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

The key constructs are based on the ability to interact by sending and receiving channel names

$x, y, z \in \mathcal{N}$	<i>Names</i>
$\pi ::= \tau \mid \bar{x}(y) \mid x(y) \mid [x = y] \pi$	<i>Prefixes</i>
$M, N ::= \mathbf{0} \mid \pi.P \mid M + M$	<i>Summations</i>
$P, Q, R ::= M \mid P_1 \mid P_2 \mid \nu x. P \mid !P$	<i>Processes</i>

π -calculus Syntax


The π -calculus is a minimal formalism that attempts to capture the essence message-passing concurrency

The key constructs are based on the ability to interact by sending and receiving channel names

$x, y, z \in \mathcal{N}$	<i>Names</i>
$\pi ::= \tau \mid \bar{x}(y) \mid x(y) \mid [x = y] \pi$	<i>Prefixes</i>
$M, N ::= \mathbf{0} \mid \pi.P \mid M + M$	<i>Summations</i>
$P, Q, R ::= M \mid P_1 \mid P_2 \mid \nu x. P \mid !P$	<i>Processes</i>

In examples, we will often appreciate $\pi.0$ as π

Reaction



Reaction

$$\overline{\tau.P + M} \rightarrow P \quad \text{R-Tau}$$

$$\overline{(\bar{x}\langle y \rangle.P_1 + M_1) \mid (x(z).P_2 + M_2)} \rightarrow P_1 \mid P_2\{y/z\} \quad \text{R-React}$$

Reaction

$$\frac{}{\tau.P + M \rightarrow P} \text{ R-Tau}$$

$$\frac{}{(\bar{x}\langle y \rangle.P_1 + M_1) \mid (x(z).P_2 + M_2) \rightarrow P_1 \mid P_2\{y/z\}} \text{ R-React}$$

$$\frac{P_1 \rightarrow P'_1}{P_1 \mid P_2 \rightarrow P'_1 \mid P_2} \text{ R-Par}$$

Reaction

$$\frac{}{\tau.P + M \rightarrow P} \text{ R-Tau}$$

$$\frac{}{(\bar{x}\langle y \rangle.P_1 + M_1) \mid (x(z).P_2 + M_2) \rightarrow P_1 \mid P_2\{y/z\}} \text{ R-React}$$

$$\frac{P_1 \rightarrow P'_1}{P_1 \mid P_2 \rightarrow P'_1 \mid P_2} \text{ R-Par}$$

$$\frac{P \rightarrow P'}{\nu x. P \rightarrow \nu x. P'} \text{ R-Res}$$

Reaction

$$\frac{}{\tau.P + M \rightarrow P} \text{R-Tau}$$

$$\frac{}{(\bar{x}\langle y \rangle.P_1 + M_1) \mid (x(z).P_2 + M_2) \rightarrow P_1 \mid P_2\{y/z\}} \text{R-React}$$

$$\frac{P_1 \rightarrow P'_1}{P_1 \mid P_2 \rightarrow P'_1 \mid P_2} \text{R-Par}$$

$$\frac{P \rightarrow P'}{\nu x. P \rightarrow \nu x. P'} \text{R-Res}$$

$$\frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q} \text{R-Struct}$$

Structural Congruence

Definition (Congruence)

An equivalence relation \mathcal{S} is a *congruence* if $P \mathcal{S} Q$ implies $C[P] \mathcal{S} C[Q]$ for every context C .

Structural Congruence

Definition (Structural Congruence)

$$[x = x] \pi.P \equiv \pi.P$$

$$!P \equiv P \mid !P$$

$$M_1 + (M_2 + M_3) \equiv (M_1 + M_2) + M_3$$

$$M_1 + M_2 \equiv M_2 + M_1$$

$$P_1 \mid (P_2 \mid P_3) \equiv (P_1 \mid P_2) \mid P_3$$

$$P_1 \mid P_2 \equiv P_2 \mid P_1$$

$$M + \mathbf{0} \equiv M$$

$$P \mid \mathbf{0} \equiv P$$

$$\nu x. \nu y. P \equiv \nu y. \nu x. P$$

$$\nu x. \mathbf{0} \equiv \mathbf{0}$$

$$\nu x. P_1 \mid P_2 \equiv P_1 \mid (\nu x. P_2), \text{ if } x \notin \text{FV}(P_1)$$

Structural Congruence

Theorem (Standard Form)

Each process is structurally congruent to one of the form

$$\nu \vec{x}. (M_1 \mid \dots \mid M_j \mid !P_1 \mid \dots \mid !P_k)$$

where each P_i is also in standard form.

Structural Congruence

Theorem (Standard Form)

Each process is structurally congruent to one of the form

$$\nu \vec{x}. (M_1 \mid \dots \mid M_j \mid !P_1 \mid \dots \mid !P_k)$$

where each P_i is also in standard form.

Proof (sketch): repeatedly use α -conversion and the scope extrusion axiom: $P \mid \nu x. Q \equiv \nu x. P \mid Q$.

Example

$$a(x).\bar{b}\langle x \rangle \mid \nu z. (\bar{a}\langle z \rangle)$$

Example

$$a(x) + b(x) \mid \nu z. (\bar{a}\langle z \rangle + \bar{b}\langle z \rangle)$$

Example

$\neg x(u). \bar{x} \langle succ\ u \rangle$

Programming in the π -calculus

Just as with λ -calculus, we can encode richer data structures and computations using the π -calculus primitives.

Polyadic π -Calculus

The send and receive primitives are **monadic**—they communicate a single name over a given channel. It is often useful to be able to send several names.

Polyadic π -Calculus

The send and receive primitives are **monadic**—they communicate a single name over a given channel. It is often useful to be able to send several names.

We can try to encode **polyadic** sends and receives as follows:

$$\bar{x}\langle y_1, \dots, y_k \rangle.P \triangleq \bar{x}\langle y_1 \rangle.\dots.\bar{x}\langle y_k \rangle.P$$

$$x(z_1, \dots, z_k).P \triangleq x(z_1).\dots.\bar{x}\langle z_k \rangle.P$$

Polyadic π -Calculus

The send and receive primitives are **monadic**—they communicate a single name over a given channel. It is often useful to be able to send several names.

We can try to encode **polyadic** sends and receives as follows:

$$\bar{x}\langle y_1, \dots, y_k \rangle.P \triangleq \bar{x}\langle y_1 \rangle.\dots.\bar{x}\langle y_k \rangle.P$$

$$x(z_1, \dots, z_k).P \triangleq x(z_1).\dots.\bar{x}\langle z_k \rangle.P$$

But unfortunately this doesn't work... why?

Polyadic π -calculus

To obtain an encoding that works correctly, we can create a **fresh name** and communicate the values over that channel:

$$\bar{x}\langle y_1, \dots, y_k \rangle.P \triangleq \nu w. (\bar{x}\langle w \rangle.\bar{w}\langle y_1 \rangle.\dots.\bar{w}\langle y_k \rangle).P$$

where $w \notin \text{FV}(P)$

$$x\langle z_1, \dots, z_k \rangle.P \triangleq x(w).w\langle z_1 \rangle.\dots.\bar{w}\langle z_k \rangle.P$$

Polyadic π -calculus

To obtain an encoding that works correctly, we can create a **fresh name** and communicate the values over that channel:

$$\bar{x}\langle y_1, \dots, y_k \rangle.P \triangleq \nu w. (\bar{x}\langle w \rangle.\bar{w}\langle y_1 \rangle.\dots.\bar{w}\langle y_k \rangle).P$$

where $w \notin \text{FV}(P)$

$$x(z_1, \dots, z_k).P \triangleq x(w).w(z_1).\dots.\bar{w}\langle z_k \rangle.P$$

Using this (adequate) encoding, we will freely use polyadic sends and receives in examples.

$$\frac{}{(\bar{x}\langle \vec{y} \rangle.P_1 + M_1) \mid (\bar{x}\langle \vec{z} \rangle.P_2 + M_2) \rightarrow P_1 \mid P_2\{\vec{y}/\vec{z}\}} \text{R-PolyReact}$$

Encoding Recursion

Idea: Suppose we want to encode P where $A(\vec{x}) \triangleq P_A$.

- Pick a name a to stand for A .
- Let $(|Q|)$ stand for Q with occurrences of $A(\vec{z})$ replaced by $\bar{a}(\vec{z})$.
- Produce $\nu a. (|P| \mid !a(\vec{x}).(|P_A|))$

Example: Buffer

Consider a recursive definition of a simple buffer:

$$\begin{aligned} B(l, r) &\triangleq r(x).C\langle x, l, r \rangle \\ C(x, l, r) &\triangleq \bar{l}\langle x \rangle.B\langle l, r \rangle \end{aligned}$$

Example: Buffer

Consider a recursive definition of a simple buffer:

$$\begin{aligned} B(l, r) &\triangleq r(x).C\langle x, l, r \rangle \\ C(x, l, r) &\triangleq \bar{l}\langle x \rangle.B\langle l, r \rangle \end{aligned}$$

When encoded this becomes

$$\nu b. \nu c. (\bar{b}\langle l, r \rangle \mid !b\langle l, r \rangle.r(x).\bar{c}\langle x, l, r \rangle \mid !c\langle x, l, r \rangle.\bar{l}\langle x \rangle.\bar{b}\langle l, r \rangle)$$