



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 steps of computation may happen simultaneously.

1 IMP with Parallel Composition

A simple way to add concurrency to a language is via a parallel composition operator. For example, we can add the $c_1 \parallel c_2$ command to IMP,

$$\begin{aligned} a & ::= x \mid n \mid a_1 + a_2 \\ b & ::= \mathbf{true} \mid \mathbf{false} \mid a_1 < a_2 \\ c & ::= \mathbf{skip} \mid x := a \mid c_1; c_2 \mid \mathbf{if } b \mathbf{ then } c_1 \mathbf{ else } c_2 \mid \mathbf{while } b \mathbf{ do } c \mid c_1 \parallel c_2, \end{aligned}$$

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 :

$$\frac{\langle \sigma, c_1 \rangle \rightarrow \langle \sigma', c'_1 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c_2 \rangle} \quad \frac{\langle \sigma, c_2 \rangle \rightarrow \langle \sigma', c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c_1 \parallel c'_2 \rangle} \quad \frac{}{\langle \sigma, \mathbf{skip} \parallel \mathbf{skip} \rangle \rightarrow \langle \sigma, \mathbf{skip} \rangle}$$

Note that the rules for parallel compositions $c_1 \parallel c_2$ allow either sub-command to take a step. In particular, the two sub-commands can interleave read and write operations involving the same store. This models a simple *shared memory* form of concurrency.

2 Communicating Sequential Processes

In the 1970s, Tony Hoare (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 so on.

CSP is organized around several new constructs for communicating information between processes. We assume a fixed collection of *channels*

$$\alpha, \beta, \gamma \in \mathbf{Channels}$$

and add new constructs $\alpha ! a$ for sending and $\alpha ? x$ for receiving messages on channels. Intuitively, if one part of a process sends a value on a channel and another process is waiting to receive, the two will interact and the value will be transmitted from one to the other. We also add a restriction operation $c \setminus \alpha$ that allows a process c to hide a channel α . This is useful for encoding internal communication within a process. Finally, we add *guards* to the language. These are based on a language construct due to Dijkstra. Intuitively, a guarded command $b \rightarrow c$ can "fire" to c if b is true and a composition of guards $g_1 \parallel g_2$ can non-deterministically "fire" to any of the commands in g_1 or g_2 whose boolean guard evaluates to true.

2.1 Syntax

The syntax of CSP is as follows. (This presentation is a slight simplification of the language given in Winskel, Chapter 14.)

$b ::= \dots$	boolean expressions
$a ::= \dots$	arithmetic expressions
$c ::= \mathbf{skip}$	skip
$x := a$	assignment
$\alpha?x$	receive
$\alpha!a$	send
$c \setminus \alpha$	restriction
$c_1 \parallel c_2$	parallel composition
$c_1; c_2$	sequential composition
if g fi	guarded conditional
do g od	guarded loop
$g ::= b \rightarrow c$	command guard
$b \wedge \alpha?x \rightarrow c$	receive guard
$b \wedge \alpha!a \rightarrow c$	send guard
$g_1 \parallel g_2$	guard composition

To ensure that all communication between processes happens via sends and receives and not via shared memory we assume a simple well-formedness condition: in every parallel composition $c_1 \parallel c_2$, the set of locations referenced in c_1 and c_2 must be disjoint.

2.2 Operational Semantics for Commands

We formalize the operational semantics of CSP using a small-step transition relation between configurations of the form $\langle \sigma, c \rangle$. To keep track of the sends and receives performed by processes, we annotate the arrow with a label λ , writing $\langle \sigma, c \rangle \xrightarrow{\lambda} \langle \sigma', c' \rangle$ to indicate that the transition has an observable event λ . The syntax of labels λ is given by the following grammar:

$$\lambda ::= \epsilon \mid \alpha!n \mid \alpha?n$$

By convention we omit the label if it is ϵ , writing $\langle \sigma, c \rangle \rightarrow \langle \sigma', c' \rangle$ instead of $\langle \sigma, c \rangle \xrightarrow{\epsilon} \langle \sigma', c' \rangle$.

The operational semantics rules are given by the axioms and rules below. The most important rules are for the parallel composition of a send and receive, which allows two processes composed in parallel to communicate a value over a channel. Some of the rules use the standard large-step semantics for arithmetic and boolean expressions in IMP. Others use an auxiliary relation that captures the small-step semantics for guards, which is defined below.

$$\frac{\langle \sigma, a \rangle \Downarrow n}{\langle \sigma, x := a \rangle \rightarrow \langle \sigma[x \mapsto n], \mathbf{skip} \rangle}$$

$$\frac{}{\langle \sigma, \alpha?x \rangle \xrightarrow{\alpha?n} \langle \sigma[x \mapsto n], \mathbf{skip} \rangle} \qquad \frac{\langle \sigma, a \rangle \Downarrow n}{\langle \sigma, \alpha!a \rangle \xrightarrow{\alpha!n} \langle \sigma, \mathbf{skip} \rangle}$$

$$\begin{array}{c}
\frac{\langle \sigma, c_1 \rangle \xrightarrow{\lambda} \langle \sigma, c'_1 \rangle}{\langle \sigma, c_1; c_2 \rangle \xrightarrow{\lambda} \langle \sigma, c'_1; c_2 \rangle} \qquad \frac{}{\langle \sigma, \mathbf{skip}; c_2 \rangle \rightarrow \langle \sigma, c_2 \rangle} \\
\\
\frac{\langle \sigma, g \rangle \rightsquigarrow \langle \sigma', c \rangle}{\langle \sigma, \mathbf{if } g \mathbf{ fi} \rangle \xrightarrow{\lambda} \langle \sigma', c \rangle} \qquad \frac{\langle \sigma, g \rangle \rightsquigarrow \langle \sigma', c \rangle}{\langle \sigma, \mathbf{do } g \mathbf{ od} \rangle \xrightarrow{\lambda} \langle \sigma', c; \mathbf{do } g \mathbf{ od} \rangle} \qquad \frac{\langle \sigma, g \rangle \rightsquigarrow \mathbf{fail}}{\langle \sigma, \mathbf{do } g \mathbf{ od} \rangle \rightarrow \langle \sigma, \mathbf{skip} \rangle} \\
\\
\frac{\langle \sigma, c_1 \rangle \xrightarrow{\lambda} \langle \sigma', c'_1 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \xrightarrow{\lambda} \langle \sigma', c'_1 \parallel c_2 \rangle} \qquad \frac{\langle \sigma, c_2 \rangle \xrightarrow{\lambda} \langle \sigma', c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \xrightarrow{\lambda} \langle \sigma', c_1 \parallel c'_2 \rangle} \qquad \frac{}{\langle \sigma, \mathbf{skip} \parallel \mathbf{skip} \rangle \rightarrow \langle \sigma, \mathbf{skip} \rangle} \\
\\
\frac{\langle \sigma, c_1 \rangle \xrightarrow{\alpha?n} \langle \sigma', c'_1 \rangle \quad \langle \sigma, c_2 \rangle \xrightarrow{\alpha!n} \langle \sigma, c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c'_2 \rangle} \qquad \frac{\langle \sigma, c_1 \rangle \xrightarrow{\alpha!n} \langle \sigma, c'_1 \rangle \quad \langle \sigma, c_2 \rangle \xrightarrow{\alpha?n} \langle \sigma', c'_2 \rangle}{\langle \sigma, c_1 \parallel c_2 \rangle \rightarrow \langle \sigma', c'_1 \parallel c'_2 \rangle} \\
\\
\frac{\langle \sigma, c \rangle \xrightarrow{\lambda} \langle \sigma', c' \rangle}{\langle \sigma, c \setminus \alpha \rangle \xrightarrow{\lambda} \langle \sigma', c' \setminus \alpha \rangle} \quad \lambda \neq \alpha?n \text{ and } \lambda \neq \alpha!n \qquad \frac{}{\langle \sigma, \mathbf{skip} \setminus \alpha \rangle \rightarrow \langle \sigma, \mathbf{skip} \rangle}
\end{array}$$

3 Operational Semantics for Guards

The small-step operational semantics for guards, written $\langle \sigma, c \rangle \xrightarrow{\lambda} \langle \sigma', c' \rangle$ or $\langle \sigma, c \rangle \rightsquigarrow \mathbf{fail}$, is a binary relation from configurations to configurations or the special value **fail**. It non-deterministically evaluates a guarded command to a configuration $\langle \sigma, c \rangle$ if the boolean expression b in one of its sub-guards of the form $b \rightarrow c$ or $b \wedge \alpha!a \rightarrow c$ or $b \wedge \alpha?x \rightarrow c$ evaluates to **true**, and produces **fail** if no such sub-guard can be found.

$$\begin{array}{c}
\frac{\langle \sigma, b \rangle \Downarrow \mathbf{true}}{\langle \sigma, b \rightarrow c \rangle \rightsquigarrow \langle \sigma, c \rangle} \qquad \frac{\langle \sigma, b \rangle \Downarrow \mathbf{false}}{\langle \sigma, b \rightarrow c \rangle \rightsquigarrow \mathbf{fail}} \\
\\
\frac{\langle \sigma, b \rangle \Downarrow \mathbf{true}}{\langle \sigma, b \wedge \alpha?x \rightarrow c \rangle \xrightarrow{\alpha?n} \langle \sigma[x \mapsto n], c \rangle} \qquad \frac{\langle \sigma, b \rangle \Downarrow \mathbf{false}}{\langle \sigma, b \wedge \alpha?x \rightarrow c \rangle \rightsquigarrow \mathbf{fail}} \\
\\
\frac{\langle \sigma, b \rangle \Downarrow \mathbf{true} \quad \langle \sigma, a \rangle \Downarrow n}{\langle \sigma, b \wedge \alpha!a \rightarrow c \rangle \xrightarrow{\alpha!n} \langle \sigma, c \rangle} \qquad \frac{\langle \sigma, b \rangle \Downarrow \mathbf{false}}{\langle \sigma, b \wedge \alpha!a \rightarrow c \rangle \rightsquigarrow \mathbf{fail}} \\
\\
\frac{\langle \sigma, g_1 \rangle \rightsquigarrow \langle \sigma', c \rangle}{\langle \sigma, g_1 \parallel g_2 \rangle \rightsquigarrow \langle \sigma', c \rangle} \qquad \frac{\langle \sigma, g_2 \rangle \rightsquigarrow \langle \sigma', c \rangle}{\langle \sigma, g_1 \parallel g_2 \rangle \rightsquigarrow \langle \sigma', c \rangle} \qquad \frac{\langle \sigma, g_1 \rangle \rightsquigarrow \mathbf{fail} \quad \langle \sigma, g_2 \rangle \rightsquigarrow \mathbf{fail}}{\langle \sigma, g_1 \parallel g_2 \rangle \rightsquigarrow \mathbf{fail}}
\end{array}$$

3.1 Examples

To get a taste for programming in CSP, let us consider a few examples.

1. Here is a simple process that computes the max of two variables and sends it on a channel m :

$$\mathbf{if} (x \leq y \rightarrow m!y) \parallel (y \leq x \rightarrow m!x) \mathbf{fi}$$

Note that the two tests, which overlap, could be evaluated concurrently on an appropriate machine.

2. Now consider a process that repeatedly receives a value on α and transmits it on β :

$$\mathbf{do} (\mathbf{true} \wedge \alpha?x \rightarrow \beta!x) \mathbf{od}$$

This implements a buffer of size one. This is useful as communication in CSP is synchronous. Using this buffer, a process that needs to send a value on α can continue without waiting for another process to receive it.

3. The next example illustrates a process that implements a buffer of size two:

$$(\mathbf{do} (\mathbf{true} \wedge \alpha?x \rightarrow \beta!x) \mathbf{od} \parallel \mathbf{do} (\mathbf{true} \wedge \beta?y \rightarrow \gamma!y) \mathbf{od}) \setminus \beta$$

Note the use of restriction to hide the internal channel β .

4. Next consider the following two processes:

$$\mathbf{if} (\mathbf{true} \wedge \alpha?x \rightarrow c_1) \parallel (\mathbf{true} \wedge \beta?x \rightarrow c_2) \mathbf{fi}$$

$$\mathbf{if} (\mathbf{true} \rightarrow \alpha?x; c_1) \parallel (\mathbf{true} \rightarrow \beta?x; c_2) \mathbf{fi}$$

These processes appear similar but they behave quite differently. Why?

5. Here is a process that merges two channels α and β onto a single channel γ :

$$\mathbf{do} (\mathbf{true} \wedge \alpha?x \rightarrow \gamma!x) \parallel (\mathbf{true} \wedge \beta?x \rightarrow \gamma!x) \mathbf{od}$$

6. This process filters a channel α , only preserving the positive values:

$$\begin{aligned} &\mathbf{do} \\ &\quad \mathbf{true} \wedge \alpha?x \rightarrow \\ &\quad \quad \mathbf{if} (x > 0 \rightarrow \beta!x) \parallel (x \leq 0 \rightarrow \mathbf{skip}) \mathbf{fi} \\ &\mathbf{od} \end{aligned}$$

7. Here is a process that implements a division operator:

$$\begin{aligned} &\mathbf{do} \\ &\quad \mathbf{true} \wedge in?x \rightarrow \\ &\quad \mathbf{true} \wedge in?y \rightarrow \\ &\quad \quad q := 0; \\ &\quad \quad r := x; \\ &\quad \quad \mathbf{do} (r \geq y \rightarrow \\ &\quad \quad \quad r := r - y; \\ &\quad \quad \quad q := q + 1) \\ &\quad \quad \mathbf{od}; \\ &\quad \quad out!q; \\ &\quad \quad out!r \\ &\mathbf{od} \end{aligned}$$

8. What does this process do? (Hint: it's like the buffer we saw before but...)

```
do true  $\wedge$   $\alpha?x \rightarrow$   
     $(\beta!x \rightarrow \alpha?x \rightarrow \mathbf{skip})$   
od
```

9. What about this one?

```
do  
  true  $\wedge$   $\alpha?x \rightarrow$   
    do  
      true  $\wedge$   $\alpha?y \rightarrow$   
         $\beta!x \rightarrow x := y$   
    od  
od
```

10. As a final example, here is a simple encoding of Dijkstra's dining philosophers in CSP (note this is *not* a solution—it can easily deadlock!):

$$phil_i = get_i!i; get_{(i+1)\%n}!i; nom!i; put_i!i; put_{(i+1)\%n}!i$$
$$fork_i = \mathbf{do} (\mathbf{true} \wedge get_i?x \rightarrow put_i?x) \mathbf{od}$$
$$eat = \mathbf{do} (\mathbf{true} \wedge nom?x \rightarrow \mathbf{skip}) \mathbf{od}$$
$$phil_0 \parallel phil_1 \parallel phil_2 \parallel phil_3 \parallel phil_4 \parallel fork_0 \parallel fork_1 \parallel fork_2 \parallel fork_3 \parallel fork_4 \parallel eat$$