


COMP3151/9154



Foundations of Concurrency

The Calculus of Communicating Systems

Johannes Åman Pohjola
UNSW
Term 2 2022

Where we are at

Last lecture, we studied *compositional* proof techniques.

This lecture, we'll take a brief detour into the world of *process algebra*, a high level formalism for describing concurrent systems.

Many of the resources for this lecture were borrowed from Graham Hutton.

CCS

The *Calculus of Communicating Systems*:

- Is a *process algebra*, a simple formal language to describe concurrent systems.

CCS

The *Calculus of Communicating Systems*:

- Is a *process algebra*, a simple formal language to describe concurrent systems.
- Is given semantics in terms of *labelled transition systems*.
- Was developed by Turing-award winner Robin Milner in the 1980s.
- Has an abstract view of synchronization that applies well to message passing.

Why do we learn this?

This gives us a symbolic way to describe our transition diagrams, and reason about them algebraically rather than diagrammatically.

Action Prefixing

Example

CLOCK₂ = tick.tock

defines a process called **CLOCK**₂ that executes the action “tick” then the action “tock” and then terminates.



The process:

CLOCK₃ = tock.tick

has the same actions as **CLOCK**₂ but arranges them in another order.

Action Prefixing

Example

$\mathbf{CLOCK}_2 = \text{tick.tock}$

defines a process called \mathbf{CLOCK}_2 that executes the action “tick” then the action “tock” and then terminates.



The process:

$\mathbf{CLOCK}_3 = \text{tock.tick}$

has the same actions as \mathbf{CLOCK}_2 but arranges them in another order.

Definition

If a is an action and P is a process, then $x.P$ is a process that executes x before P . This brackets to the right, so:

$$x.y.z.P = x.(y.(z.P))$$

Action Prefixing

Example

$\mathbf{CLOCK}_2 = \text{tick.tock}$

defines a process called \mathbf{CLOCK}_2 that executes the action “tick” then the action “tock” and then terminates.



The process:

$\mathbf{CLOCK}_3 = \text{tock.tick}$

has the same actions as \mathbf{CLOCK}_2 but arranges them in another order.

Definition

If a is an action and P is a process, then $x.P$ is a process that executes x before P . This brackets to the right, so:

$$x.y.z.P = x.(y.(z.P))$$

Stopping

More precisely, we should write:

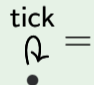
$\mathbf{CLOCK}_2 = \text{tick.tock.STOP}$

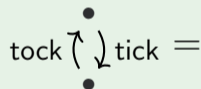
where \mathbf{STOP} is the trivial process with no transitions.

Loops

Up to now, all processes make a **finite** number of transitions and then terminate. Processes that can make a **infinite** number of transitions can be pictured by allowing **loops**:

Example (Loops)


 = the process that diverges
 executing "tick" transitions



 = the process that alternates
 "tick" and "tock" forever

We accomplish loops in CCS using **recursion**.

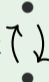
Loops

Up to now, all processes make a **finite** number of transitions and then terminate. Processes that can make a **infinite** number of transitions can be pictured by allowing **loops**:

Example (Loops)


 $\text{tick} =$ the process that diverges
executing “tick” transitions

$\mathbf{CLOCK}_4 = \text{tick}.\mathbf{CLOCK}_4$


 $\text{tock} \left(\right) \text{tick} =$ the process that alternates
“tick” and “tock” forever

$\mathbf{CLOCK}_5 = \text{tick}.\text{tock}.\mathbf{CLOCK}_5$

We accomplish loops in CCS using **recursion**.

Equality of Processes

These two processes are physically different:

$$\begin{array}{c}
 \text{tick} \\
 \curvearrowright \\
 \bullet \\
 \text{CLOCK}_4 = \text{tick.CLOCK}_4
 \end{array}$$


$$\begin{array}{c}
 \bullet \\
 \text{tick} \left(\begin{array}{l} \uparrow \\ \downarrow \end{array} \right) \text{tick} \\
 \bullet \\
 \text{CLOCK}_6 = \text{tick.tick.CLOCK}_6
 \end{array}$$

Equality of Processes

These two processes are physically different:



$$\mathbf{CLOCK}_4 = \text{tick}.\mathbf{CLOCK}_4$$



$$\mathbf{CLOCK}_6 = \text{tick}.\text{tick}.\mathbf{CLOCK}_6$$

But they both have the same **behaviour** — an infinite sequence of “tick” transitions.

Informal Definition

We consider two process to be equal if an external observer cannot distinguish them by their actions.

We will refine this definition later.


A Simple Vending Machine

Vending Machines are very common examples for process algebra.

Example (An inflexible machine)

Suppose I define my vending machine as:

$$\mathbf{VM}_1 = \text{in}50\text{¢}.\text{outCoke}.\text{in}20\text{¢}.\text{outMars}.\mathbf{VM}_1$$




A Simple Vending Machine

Vending Machines are very common examples for process algebra.

Example (An inflexible machine)

Suppose I define my vending machine as:

$$\mathbf{VM}_1 = \text{in}50\text{¢}.\text{outCoke}.\text{in}20\text{¢}.\text{outMars}.\mathbf{VM}_1$$



This machine is not very flexible:

- It only accepts exact money.
- The customer has no choice: The machine dispenses Coke and Mars bars alternately.

Choice


To make a more flexible kind of vending machine, we need a (nondeterministic) **choice** operator.

Choice

To make a more flexible kind of vending machine, we need a (nondeterministic) **choice** operator.

Example

$$\mathbf{VM}_2 = (\text{in}50\text{¢}.\text{outCoke}.\mathbf{VM}_2) + (\text{in}20\text{¢}.\text{outMars}.\mathbf{VM}_2)$$

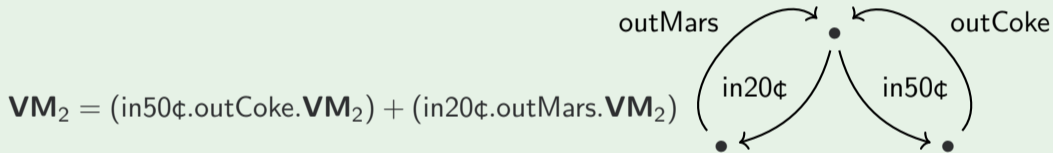


Here we have a process \mathbf{VM}_2 that repeatedly **either** inputs 50¢ and outputs a coke, **or** inputs 20¢ and outputs a mars bar.

Choice

To make a more flexible kind of vending machine, we need a (nondeterministic) **choice** operator.

Example



Here we have a process \mathbf{VM}_2 that repeatedly **either** inputs 50¢ and outputs a coke, **or** inputs 20¢ and outputs a mars bar.

Definition

If P and Q are processes then $P + Q$ is a process which can either behave as the process P or the process Q .

Choice Equalities

Observe that we have the following identities about choice:

$$P + (Q + R) = (P + Q) + R \quad (\text{associativity})$$

Choice Equalities

Observe that we have the following identities about choice:

$$P + (Q + R) = (P + Q) + R \quad (\text{associativity})$$

$$P + Q = Q + P \quad (\text{commutativity})$$

Choice Equalities

Observe that we have the following identities about choice:

$$P + (Q + R) = (P + Q) + R \quad (\text{associativity})$$

$$P + Q = Q + P \quad (\text{commutativity})$$

$$P + \mathbf{STOP} = P \quad (\text{neutral element})$$

Choice Equalities

Observe that we have the following identities about choice:

$$P + (Q + R) = (P + Q) + R \quad (\text{associativity})$$

$$P + Q = Q + P \quad (\text{commutativity})$$

$$P + \mathbf{STOP} = P \quad (\text{neutral element})$$

$$P + P = P \quad (\text{idempotence})$$

Choice Equalities

Observe that we have the following identities about choice:

$$P + (Q + R) = (P + Q) + R \quad (\text{associativity})$$


$$P + Q = Q + P \quad (\text{commutativity})$$

$$P + \mathbf{STOP} = P \quad (\text{neutral element})$$

$$P + P = P \quad (\text{idempotence})$$

What about the equation:

$$a.(P + Q) \stackrel{?}{=} (a.P) + (a.Q)$$





Branching Time

Example

$$VM_3 = \text{in}50\text{¢}.\text{(outCoke + outPepsi)}$$

$$VM_4 = (\text{in}50\text{¢}.\text{outCoke}) + (\text{in}50\text{¢}.\text{outPepsi})$$





Branching Time

Example

$$\mathbf{VM}_3 = \text{in}50\text{¢} \cdot (\text{outCoke} + \text{outPepsi})$$

$$\mathbf{VM}_4 = (\text{in}50\text{¢} \cdot \text{outCoke}) + (\text{in}50\text{¢} \cdot \text{outPepsi})$$



Reactive Systems



\mathbf{VM}_3 allows the **customer** to choose which drink to vend after inserting 50¢. In \mathbf{VM}_4 however, the **machine** makes the choice when the customer inserts a coin.

Branching Time

Example

$$\mathbf{VM}_3 = \text{in50}\text{¢}.\text{(outCoke} + \text{outPepsi)}$$

$$\mathbf{VM}_4 = (\text{in50}\text{¢}.\text{outCoke}) + (\text{in50}\text{¢}.\text{outPepsi})$$



Or in pictures:

Reactive Systems

\mathbf{VM}_3 allows the **customer** to choose which drink to vend after inserting 50¢. In \mathbf{VM}_4 however, the **machine** makes the choice when the customer inserts a coin.

They are **different** in this *reactive* view, but they have the **same behaviours!**

Equivalences

The equation

$$a.(P + Q) = (a.P) + (a.Q)$$

is usually not admitted for this reason.

Equivalences

The equation

$$a.(P + Q) = (a.P) + (a.Q)$$

is usually not admitted for this reason.

Exercise

It is possible to construct two processes that are equal assuming this equation but do not have the same set of behaviours (and thus do not satisfy the same LTL properties).

Equivalences

The equation

$$a.(P + Q) = (a.P) + (a.Q)$$

is usually not admitted for this reason.

Exercise

It is possible to construct two processes that are equal assuming this equation but do not have the same set of behaviours (and thus do not satisfy the same LTL properties).

If we **do** admit it, then our notion of equality is very coarse (it is called **partial trace equivalence**). This is enough if we want to prove safety properties, but **progress** is not guaranteed. Johannes: Explain why

Equivalences

The equation

$$a.(P + Q) = (a.P) + (a.Q)$$

is usually not admitted for this reason.

Exercise

It is possible to construct two processes that are equal assuming this equation but do not have the same set of behaviours (and thus do not satisfy the same LTL properties).

If we **do** admit it, then our notion of equality is very coarse (it is called **partial trace equivalence**). This is enough if we want to prove safety properties, but **progress** is not guaranteed. Johannes: Explain why

Terminology

Our notion of equality without this equation is called (strong) **bisimilarity**.

Exercises

- A clock that can stop at any time.

Exercises

- A clock that can stop at any time.
- A clock that ticks or tocks at each cycle.

Exercises

- A clock that can stop at any time.
- A clock that ticks or tocks at each cycle.
- A clock that ticks each cycle or tocks each cycle.

Exercises

- A clock that can stop at any time.
- A clock that ticks or tocks at each cycle.
- A clock that ticks each cycle or tocks each cycle.
- A vending machine for Mars and Coke that gives change.

Parallel Composition

Definition

If P and Q are processes then $P \mid Q$ is the parallel composition of their processes — i.e. the non-deterministic interleaving of their actions.


Parallel Composition

Definition

If P and Q are processes then $P \mid Q$ is the parallel composition of their processes — i.e. the non-deterministic interleaving of their actions.

Example (Clocks)

ACLOCK = tick.beep | tock



CCLOCK = TICLK|TOCLK

TICLK = tick.TICLK

TOCLK = tock.TOCLK

Exercise: Express these processes without parallel composition.

Synchronization

In CCS, every action a has an opposing *coaction* \bar{a} (and $\bar{\bar{a}} = a$):

Actions: tick tock in50¢ outCoke ...

Coactions: $\overline{\text{tick}}$ $\overline{\text{tock}}$ $\overline{\text{in50¢}}$ $\overline{\text{outCoke}}$...

Synchronization

In CCS, every action a has an opposing *coaction* \bar{a} (and $\bar{\bar{a}} = a$):

Actions: tick tock in50¢ outCoke ...

Coactions: $\overline{\text{tick}}$ $\overline{\text{tock}}$ $\overline{\text{in50¢}}$ $\overline{\text{outCoke}}$...

It is a convention to think of an action as an **output** event and a coaction as an **input** event. If a system can execute both an action and its coaction, it **may** execute them both simultaneously by taking an **internal transition** marked by the special action τ .

Synchronization

In CCS, every action a has an opposing *coaction* \bar{a} (and $\bar{\bar{a}} = a$):

Actions: tick tock in50¢ outCoke ...

Coactions: $\overline{\text{tick}}$ $\overline{\text{tock}}$ $\overline{\text{in50¢}}$ $\overline{\text{outCoke}}$...


It is a convention to think of an action as an **output** event and a coaction as an **input** event. If a system can execute both an action and its coaction, it **may** execute them both simultaneously by taking an **internal transition** marked by the special action τ .

Example (Relay Race)

$\text{RACE} = \text{RUN}_1 \mid \text{RUN}_2$

$\text{RUN}_1 = \text{start.baton}$

$\text{RUN}_2 = \overline{\text{baton}}.\text{finish}$



Expansion Theorem

Let P and Q be processes. By expanding recursive definitions and using our existing equations for choice we can express P and Q as n -ary choices of action prefixes:

$$P = \sum_{i \in I} \alpha_i. P_i \text{ and } Q = \sum_{j \in J} \beta_j. Q_j.$$

Expansion Theorem

Let P and Q be processes. By expanding recursive definitions and using our existing equations for choice we can express P and Q as n -ary choices of action prefixes:

$$P = \sum_{i \in I} \alpha_i. P_i \text{ and } Q = \sum_{j \in J} \beta_j. Q_j.$$

Then, the parallel composition can be expressed as follows:

$$P \mid Q = \sum_{i \in I} \alpha_i.(P_i \mid Q) + \sum_{j \in J} \beta_j.(P \mid Q_j) + \sum_{i \in I, j \in J, \alpha_i = \bar{\beta}_j} \tau.(P_i \mid Q_j).$$

Expansion Theorem

Let P and Q be processes. By expanding recursive definitions and using our existing equations for choice we can express P and Q as n -ary choices of action prefixes:

$$P = \sum_{i \in I} \alpha_i. P_i \text{ and } Q = \sum_{j \in J} \beta_j. Q_j.$$

Then, the parallel composition can be expressed as follows:

$$P \mid Q = \sum_{i \in I} \alpha_i.(P_i \mid Q) + \sum_{j \in J} \beta_j.(P \mid Q_j) + \sum_{i \in I, j \in J, \alpha_i = \bar{\beta}_j} \tau.(P_i \mid Q_j).$$

From this, many useful equations are derivable:

$$\begin{aligned} P \mid Q &= Q \mid P \\ P \mid (Q \mid R) &= (P \mid Q) \mid R \\ P \mid \mathbf{STOP} &= P \end{aligned}$$

Restriction

We wish a way to say “these are all the processes that there are”, in other words, to **force** synchronization to happen and not allow certain actions to be taken alone.

Restriction

We wish a way to say “these are all the processes that there are”, in other words, to **force** synchronization to happen and not allow certain actions to be taken alone.

Definition

If P is a process and a is an action (not τ), then $P \setminus a$ is the same as the process P **except** that the actions a and \bar{a} may not be executed. We have

$$(a.P) \setminus b = a.(P \setminus b) \quad \text{if } a \notin \{b, \bar{b}\}$$

Restriction

We wish a way to say “these are all the processes that there are”, in other words, to **force** synchronization to happen and not allow certain actions to be taken alone.

Definition

If P is a process and a is an action (not τ), then $P \setminus a$ is the same as the process P **except** that the actions a and \bar{a} may not be executed. We have


$$(a.P) \setminus b = a.(P \setminus b) \quad \text{if } a \notin \{b, \bar{b}\}$$

Example (Relay Race)

$$\text{RACE} = (\text{RUN}_1 \mid \text{RUN}_2) \setminus \text{baton}$$

$$\text{RUN}_1 = \text{start.baton}$$

$$\text{RUN}_2 = \overline{\text{baton}}.\text{finish}$$



Another Example

A man that eats every time a clock ticks:

$$\mathbf{CLOCK}_4 = \text{tick}.\mathbf{CLOCK}_4$$

$$\mathbf{MAN} = \overline{\text{tick}}.\text{eat}.\mathbf{MAN}$$

$$\mathbf{EXAMPLE} = (\mathbf{MAN} \mid \mathbf{CLOCK}_4) \setminus \text{tick}$$

Another Example


A man that eats every time a clock ticks:

$$\mathbf{CLOCK}_4 = \text{tick}.\mathbf{CLOCK}_4$$

$$\mathbf{MAN} = \overline{\text{tick}}.\text{eat}.\mathbf{MAN}$$

$$\mathbf{EXAMPLE} = (\mathbf{MAN} \mid \mathbf{CLOCK}_4) \setminus \text{tick}$$

After deriving the picture, we get:




Semantics

Up until now, our semantics were given informally in terms of pictures. Now we will formalise our semantic intuitions.

Semantics

Up until now, our semantics were given informally in terms of pictures. Now we will formalise our semantic intuitions.

Our set of locations in our labelled transition system will be the **set of all CCS processes**. Locations can now be labelled with what process they are:



We will now define what transitions exist in our LTS by means of a set of *inference rules*. This technique is called *operational semantics*.

Inference Rules

In logic we often write:

$$\frac{A_1 \quad A_2 \quad \dots \quad A_n}{C}$$

To indicate that C can be proved by proving all assumptions A_1 through A_n .
For example, the classical logical rule of **modus ponens** is written as follows:

$$\frac{A \Rightarrow B \quad A}{B} \text{ MODUS PONENS}$$

Operational Semantics

Operational Semantics

$$\frac{}{a.P \xrightarrow{a} P} \text{ACT} \quad \frac{P \xrightarrow{a} P'}{P + Q \xrightarrow{a} P'} \text{CHOICE}_1 \quad \frac{Q \xrightarrow{a} Q'}{P + Q \xrightarrow{a} Q'} \text{CHOICE}_2$$

$$\frac{P \xrightarrow{a} P'}{P | Q \xrightarrow{a} P' | Q} \text{PAR}_1 \quad \frac{Q \xrightarrow{a} Q'}{P | Q \xrightarrow{a} P | Q'} \text{PAR}_2 \quad \frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\bar{a}} Q'}{P | Q \xrightarrow{\tau} P' | Q'} \text{SYNC}$$

$$\frac{P \xrightarrow{a} P' \quad a \notin \{b, \bar{b}\}}{P \setminus b \xrightarrow{a} P' \setminus b} \text{RESTRICT}$$

Operational Semantics


$$\begin{array}{c}
 \frac{}{a.P \xrightarrow{a} P} \text{ACT} \quad \frac{P \xrightarrow{a} P'}{P + Q \xrightarrow{a} P'} \text{CHOICE}_1 \quad \frac{Q \xrightarrow{a} Q'}{P + Q \xrightarrow{a} Q'} \text{CHOICE}_2 \\
 \\
 \frac{P \xrightarrow{a} P'}{P | Q \xrightarrow{a} P' | Q} \text{PAR}_1 \quad \frac{Q \xrightarrow{a} Q'}{P | Q \xrightarrow{a} P | Q'} \text{PAR}_2 \quad \frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\bar{a}} Q'}{P | Q \xrightarrow{\tau} P' | Q'} \text{SYNC} \\
 \\
 \frac{P \xrightarrow{a} P' \quad a \notin \{b, \bar{b}\}}{P \setminus b \xrightarrow{a} P' \setminus b} \text{RESTRICT}
 \end{array}$$

Bisimulation Equivalence

Two processes (or locations) P and Q are **bisimilar** iff they can do the same actions and those actions themselves lead to bisimilar processes. All of our previous equalities can be proven by induction on the semantics here.

Proof Trees

The advantages of this rule presentation is that they can be “stacked” to give a neat tree like derivation of proofs.



Exercise: Show $((a.P) + Q) \mid \bar{a}.R \xrightarrow{\tau} P \mid R$

Value Passing

We add synchronous channels into CCS by letting actions take **parameters**.

Actions: $a(3)$ $c(15)$ $x(\textit{True})$...

Coactions: $\bar{a}(x)$ $\bar{c}(y)$ $\bar{c}(z)$...

Value Passing

We add synchronous channels into CCS by letting actions take **parameters**.

Actions: $a(3)$ $c(15)$ $x(\textit{True})$...

Coactions: $\bar{a}(x)$ $\bar{c}(y)$ $\bar{c}(z)$...

The parameter of an action is the value to be sent, and the parameter of a coaction is the **variable** in which the received value is stored.

Value Passing

We add synchronous channels into CCS by letting actions take **parameters**.

Actions: $a(3)$ $c(15)$ $x(\text{True})$...

Coactions: $\bar{a}(x)$ $\bar{c}(y)$ $\bar{c}(z)$...

The parameter of an action is the value to be sent, and the parameter of a coaction is the **variable** in which the received value is stored.

Example (Buffers)

A one-cell sized buffer is implemented as:

$$\mathbf{BUFF} = \bar{\text{in}}(x).\text{out}(x).\mathbf{BUFF}$$

Larger buffers can be made by stitching multiple **BUFF** processes together! This is one (overkill) way to model asynchronous communication in CCS.

Guards

Rather than add **if** statements, we add the notion of a *guard*:

Guards

Rather than add **if** statements, we add the notion of a *guard*:

Definition

If P is a value-passing CCS process and φ is a formula about the variables in scope, then $[\varphi]P$ is a process that executes just like P if φ holds for the current state, and like **STOP** otherwise.

Guards

Rather than add **if** statements, we add the notion of a *guard*:

Definition

If P is a value-passing CCS process and φ is a formula about the variables in scope, then $[\varphi]P$ is a process that executes just like P if φ holds for the current state, and like **STOP** otherwise.

We can define an **if** statement like so:

$$\mathbf{if} \varphi \mathbf{ then } P \mathbf{ else } Q \quad \equiv \quad ([\varphi].P) + ([\neg\varphi].Q)$$

Assignment

Most process algebras have no notion of state. Some presentations of value passing CCS also include assignment to update variables in the state:

Assignment

Most process algebras have no notion of state. Some presentations of value passing CCS also include assignment to update variables in the state:

Definition

If P is a process and x is a variable in the state, and e is an expression, then $\llbracket x := e \rrbracket P$ is the same as P except that it first updates the variable x to have the value e before making the transition.

Assignment

Most process algebras have no notion of state. Some presentations of value passing CCS also include assignment to update variables in the state:

Definition

If P is a process and x is a variable in the state, and e is an expression, then $\llbracket x := e \rrbracket P$ is the same as P except that it first updates the variable x to have the value e before making the transition.

With this, our value-passing CCS is now just as expressive as Ben-Ari's pseudocode. Moreover, the connection between CCS and transition diagrams is formalised, enabling us to reason symbolically about processes rather than semantically.

Process Algebra

This was an example of a *process algebra*. There are many such algebras and they have been very influential on the design of concurrent programming languages.

Process Algebra

This was an example of a *process algebra*. There are many such algebras and they have been very influential on the design of concurrent programming languages.

Other process algebras include:

- The *Algebra of Communicating Processes* (Bergstra and Klop, 1982) which distinguishes between deadlock and termination.
- The *Communicating Sequential Processes* formalism (Hoare, 1978) with a more refined treatment of nondeterminism.
- The *π -calculus* (Milner et al. 1992), a derivative of CCS that allows for first class channels and processes.

Process Algebra

This was an example of a *process algebra*. There are many such algebras and they have been very influential on the design of concurrent programming languages.

Other process algebras include:

- The *Algebra of Communicating Processes* (Bergstra and Klop, 1982) which distinguishes between deadlock and termination.
- The *Communicating Sequential Processes* formalism (Hoare, 1978) with a more refined treatment of nondeterminism.
- The *π -calculus* (Milner et al. 1992), a derivative of CCS that allows for first class channels and processes.

There are dozens of equivalences other than strong bisimulation that are useful for various scenarios.

