

Concurrency Theory

Rob van Glabbeek

University of Edinburgh

University of New South Wales, Sydney, Australia

Stanford University

July 2025

4 hour course on concurrency theory

- Modelling Distributed systems
 - Labelled Transition Systems (Monday)
 - Petri nets (Thursday, if time allows, on whiteboard)
 - Process algebra – CCSP (Monday)
(Denotational semantics on blackboard)
- When does an implementation meet a specification?
 - Semantic equivalences and preorders
(Monday, after the break, on whiteboard)
- Does a system have a property?
 - Safety and liveness properties (Monday & Tuesday)
 - Temporal logic (Tuesday)
 - Progress, Justness and Fairness assumptions (Tuesday & Thursday)
- Session types (Thursday)
- Mutual exclusion (Thursday)

30-hour course on this subject, with recorded lectures:

<https://cgi.cse.unsw.edu.au/~rvg/6752>

Specification

what we want

Implementation

what we made

Object
or
Expression
in
mathematical Model
or
Specification Language

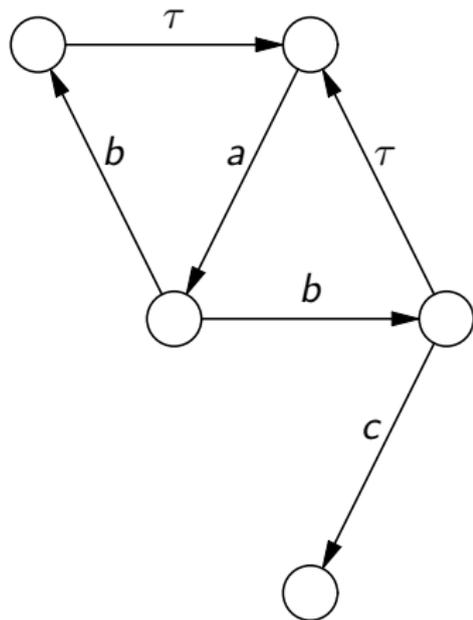
Object
or
Expression
in
math. Model
or
Language

?

Verification

Semantic Equivalence or preorder

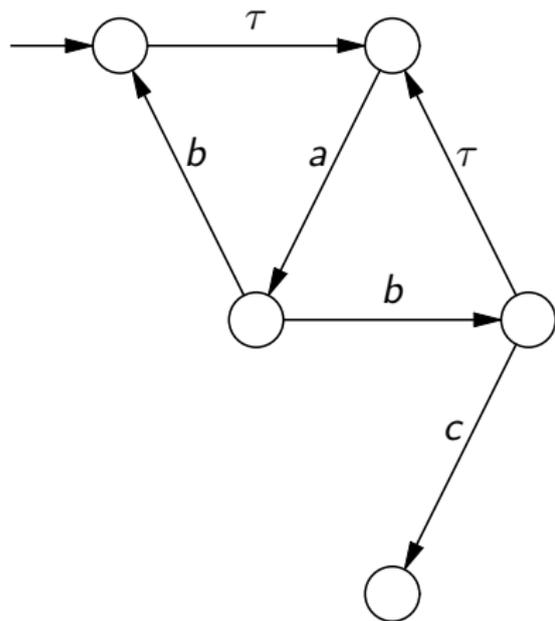
Labelled Transition Systems



- τ internal action
- a insert coin in slot
- b dispense pretzel
- c close down machine

A *labelled transition system*

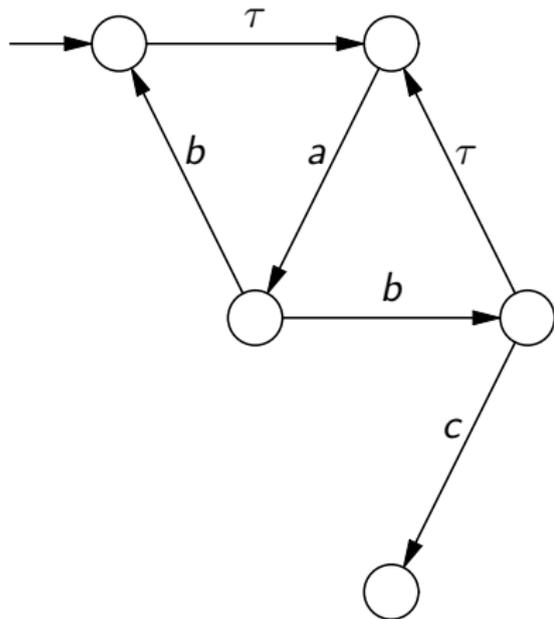
Labelled Transition Systems



- τ internal action
- a insert coin in slot
- b dispense pretzel
- c close down machine

A *process graph*

Labelled Transition Systems



- τ internal action
- a insert coin in slot
- b dispense pretzel
- c close down machine

A *process graph*

(completed) traces:
 $(ab)^\infty \cup ab(ab)^*c.$

Labelled Transition Systems

Let Act be a set of *actions*.

A *Labelled Transition System* over Act is tuple (S, \rightarrow) with

- S a set (of *states*), and
- $\rightarrow \subseteq S \times Act \times S$, the *transition relation*.

Process graphs

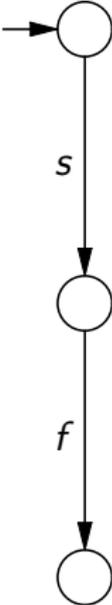
Let Act be a set of *actions*.

A *Process graph* over Act is tuple (S, \rightarrow, I) with

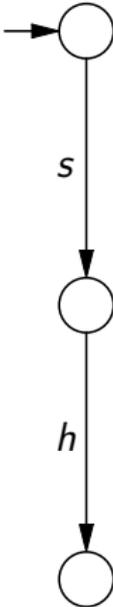
- S a set (of *states*)
- $\rightarrow \subseteq S \times Act \times S$, the *transition relation*, and
- $I \in S$ the initial state.

Relay Race

Specification



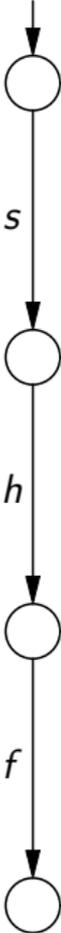
s = start
h = handover
f = finish



Implementation



=



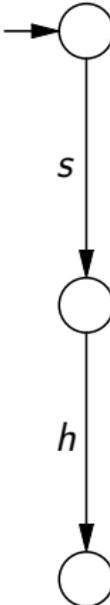
Relay Race

Specification



s = start
h = handover
f = finish

Abstraction

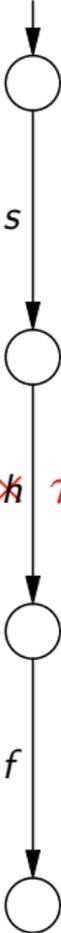


Implementation



\parallel_h

$=$



The process algebra CCSP

$P, Q ::= 0 \mid \alpha.P \mid P + Q \mid P \parallel_S Q \mid \tau_I(P) \mid \text{Relabel} \mid \text{Recursion}$

with $\alpha \in \text{Act} := A \uplus \{\tau\}$, $S, I \subseteq A$.

The process algebra CCSP

$P, Q ::= 0 \mid \alpha.P \mid P + Q \mid P \parallel_S Q \mid \tau_I(P) \mid \text{Relabel} \mid \text{Recursion}$

with $\alpha \in \text{Act} := A \uplus \{\tau\}$, $S, I \subseteq A$.

s.h.0

The process algebra CCSP

$P, Q ::= 0 \mid \alpha.P \mid P + Q \mid P \parallel_S Q \mid \tau_I(P) \mid \text{Relabel} \mid \text{Recursion}$

with $\alpha \in \text{Act} := A \uplus \{\tau\}$, $S, I \subseteq A$.

s.h.0

coin.(chocolate.0 + pretzel.0)

VM := coin.(chocolate.VM + pretzel.VM)

The process algebra CCSP

$P, Q ::= 0 \mid \alpha.P \mid P + Q \mid P \parallel_S Q \mid \tau_I(P) \mid \text{Relabel} \mid \text{Recursion}$

with $\alpha \in \text{Act} := A \uplus \{\tau\}$, $S, I \subseteq A$.

s.h.0

coin.(chocolate.0 + pretzel.0)

VM := coin.(chocolate.VM + pretzel.VM)

$\tau_{\{h\}}(s.h \parallel_{\{h\}} h.f)$

Operational semantics of CCSP

$$\begin{array}{c} \alpha.P \xrightarrow{\alpha} P \\ \frac{P \xrightarrow{\alpha} P'}{P + Q \xrightarrow{\alpha} P'} \quad \frac{Q \xrightarrow{\alpha} Q'}{P + Q \xrightarrow{\alpha} Q'} \quad \frac{P \xrightarrow{\alpha} P'}{\mathcal{R}(P) \xrightarrow{\beta} \mathcal{R}(P')} \quad \left(\begin{array}{l} \alpha = \beta = \tau \vee \\ (\alpha, \beta) \in \mathcal{R} \end{array} \right) \\ \frac{P \xrightarrow{\alpha} P'}{P \parallel_S Q \xrightarrow{\alpha} P' \parallel_S Q} \quad (\alpha \notin S) \quad \frac{P \xrightarrow{a} P' \quad Q \xrightarrow{a} Q'}{P \parallel_S Q \xrightarrow{a} P' \parallel_S Q'} \quad (a \in S) \quad \frac{Q \xrightarrow{\alpha} Q'}{P \parallel_S Q \xrightarrow{\alpha} P \parallel_S Q'} \quad (\alpha \notin S) \\ \frac{P \xrightarrow{\alpha} P'}{\tau_I(P) \xrightarrow{\alpha} \tau_I(P')} \quad (\alpha \notin I) \quad \frac{P \xrightarrow{a} P'}{\tau_I(P) \xrightarrow{\tau} \tau_I(P')} \quad (a \in I) \quad \frac{P \xrightarrow{\alpha} P'}{X \xrightarrow{\alpha} P'} \quad (X := P) \end{array}$$

Homework

Consider a process with the following set of completed traces:
 $\{bad, abd, adb, acd, adc\}$. Thus, in each run, the actions a, d and one out of c or b occur once each, subject to the following restrictions:

- b and c can never occur in the same trace.
- if c happens, it must be after a, and
- a happens before d.

Give a CCSP expression for this process that uses at most one $+$.

Compositionality

$$P_1 \parallel \cdots \parallel P_n$$

n parallel components of k states each: n^k states total.

state explosion problem

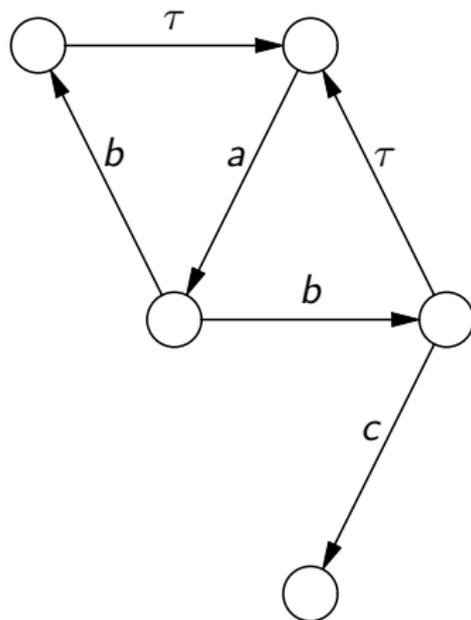
If we replace P_i by an equivalent component Q_i , for $i = 1, \dots, n$ then we want the new system $Q_1 \parallel \cdots \parallel Q_n$ to be equivalent to the old system $P_1 \parallel \cdots \parallel P_n$.

This is called *compositionality*.

Next to abstraction, this is the most important tool in combatting the state explosion problem.

Safety and Liveness

- τ internal action
- a insert coin in slot
- b dispense pretzel
- c close down machine



Safety:

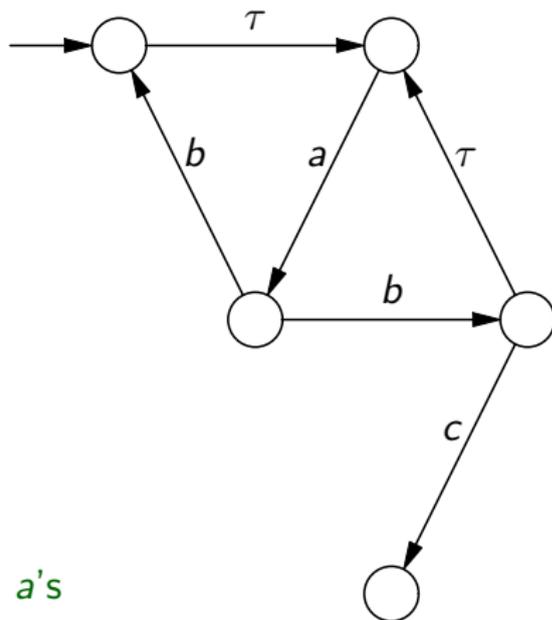
nothing bad will ever happen

Liveness:

something good will ever happen eventually

Safety and Liveness

- τ internal action
- a insert coin in slot
- b dispense pretzel
- c close down machine



Safety:

nothing bad will ever happen

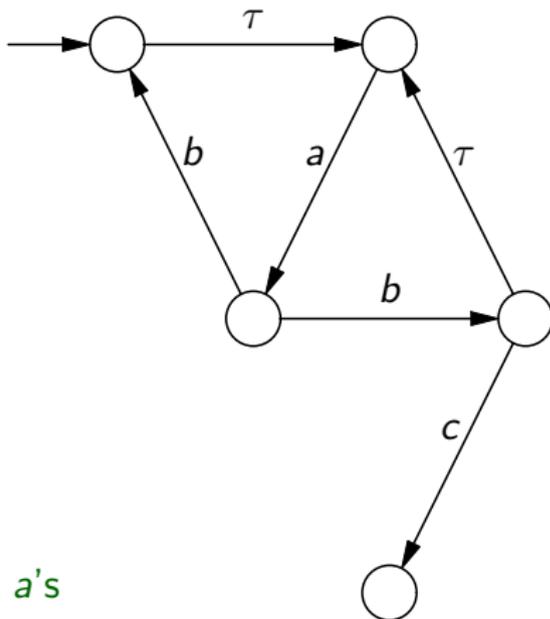
no run contains more b 's than a 's

Liveness:

something good will ever happen eventually

Safety and Liveness

- τ internal action
- a insert coin in slot
- b dispense pretzel
- c close down machine



Safety:

nothing bad will ever happen

no run contains more b 's than a 's

Liveness:

something good will ever happen eventually

each a will eventually be followed by a b

VM = a.VM

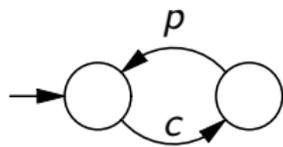
Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).

$$VM = c.p.VM$$

Pretzels

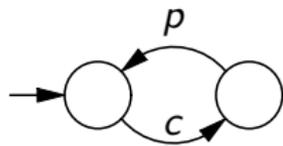
Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



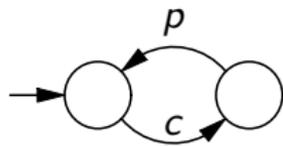
$$VM = c.p.VM$$

properties of systems.

whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



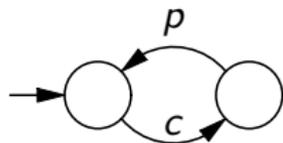
$$VM = c.p.VM$$

(desired) properties of systems.

whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



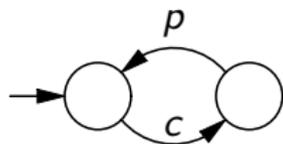
$$VM = c.p.VM$$

Temporal Logic is a formalism for specifying (desired) properties of systems.

whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



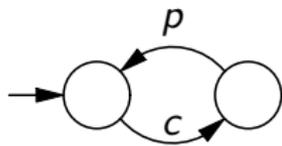
$$VM = c.p.VM$$

Temporal Logic is a formalism for specifying (desired) properties of systems. It can tell what should happen, and in which order, but not when.

whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



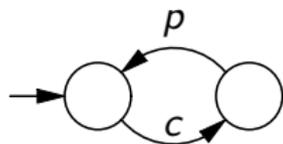
$$VM = c.p.VM$$

Linear-time Temporal Logic (LTL) is a formalism for specifying (desired) properties of systems. It can tell what should happen, and in which order, but not when. [Pnueli 1977]

whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$

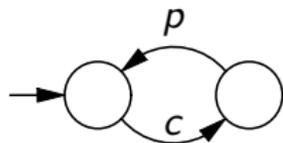
Linear-time Temporal Logic (LTL) is a formalism for specifying (desired) properties of systems. It can tell what should happen, and in which order, but not when. [Pnueli 1977]

In LTL we have $(c \Rightarrow p)$. \leftarrow formula

whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$

Linear-time Temporal Logic (LTL) is a formalism for specifying (desired) properties of systems. It can tell what should happen, and in which order, but not when. [Pnueli 1977]

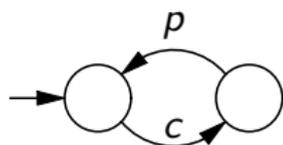
In LTL we have $(c \Rightarrow \mathbf{F}p)$. ← a temporal formula

whenever a coin is inserted, eventually a pretzel is produced.

F (or \diamond) means **eventually**

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$

Linear-time Temporal Logic (LTL) is a formalism for specifying (desired) properties of systems. It can tell what should happen, and in which order, but not when. [Pnueli 1977]

In LTL we have $\mathbf{G}(c \Rightarrow \mathbf{F}p)$. \leftarrow a temporal formula

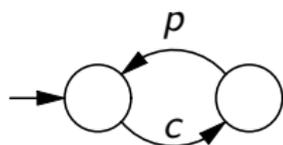
whenever a coin is inserted, eventually a pretzel is produced.

F (or \diamond) means **eventually**

G (or \square) means **globally**, in each reachable state

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$

Linear-time Temporal Logic (LTL) is a formalism for specifying (desired) properties of systems. It can tell what should happen, and in which order, but not when. [Pnueli 1977]

In LTL we have $VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$. \leftarrow a temporal judgement

whenever a coin is inserted, eventually a pretzel is produced.

F (or \diamond) means **eventually**

G (or \square) means **globally**, in each reachable state

\models means **satisfies**.

Kripke structures

Let AP be a set of *atomic predicates*.

A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with

- S a set (of *states*),
- $\rightarrow \subseteq S \times S$, the *transition relation*, and
- $\models \subseteq S \times AP$.

$s \models p$ says that predicate $p \in AP$ *holds* in state $s \in S$.

Kripke structures

Let AP be a set of *atomic predicates*.

A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with

- S a set (of *states*),
- $\rightarrow \subseteq S \times S$, the *transition relation*, and
- $\models \subseteq S \times AP$.

$s \models p$ says that predicate $p \in AP$ *holds* in state $s \in S$.

A *path* in a Kripke structure is a nonempty finite or infinite sequence s_0, s_1, \dots of states that follows the transition relation.

Kripke structures

Let AP be a set of *atomic predicates*.

A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with

- S a set (of *states*),
- $\rightarrow \subseteq S \times S$, the *transition relation*, and
- $\models \subseteq S \times AP$.

$s \models p$ says that predicate $p \in AP$ *holds* in state $s \in S$.

A *path* in a Kripke structure is a nonempty finite or infinite sequence s_0, s_1, \dots of states that follows the transition relation.

If ρ is a finite prefix of π , then $\pi \upharpoonright \rho$ denotes the suffix of π that remains after removing the prefix ρ , but not the last state of ρ .

Kripke structures

Let AP be a set of *atomic predicates*.

A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with

- S a set (of states),
- $\rightarrow \subseteq S \times S$, the *transition relation*, and
- $\models \subseteq S \times AP$.

$s \models p$ says that predicate $p \in AP$ *holds* in state $s \in S$.

A *path* in a Kripke structure is a nonempty finite or infinite sequence s_0, s_1, \dots of states that follows the transition relation.

If ρ is a finite prefix of π , then $\pi \upharpoonright \rho$ denotes the suffix of π that remains after removing the prefix ρ , but not the last state of ρ .

The *length* of a path π , denoted $|\pi| \in \mathbb{N} \cup \{\infty\}$, is the number of transitions in π ; for instance $|s_0 s_1 s_2 s_3| = 3$.

Kripke structures

Let AP be a set of *atomic predicates*.

A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with

- S a set (of states),
- $\rightarrow \subseteq S \times S$, the *transition relation*, and
- $\models \subseteq S \times AP$.

$s \models p$ says that predicate $p \in AP$ *holds* in state $s \in S$.

A *path* in a Kripke structure is a nonempty finite or infinite sequence s_0, s_1, \dots of states that follows the transition relation.

If ρ is a finite prefix of π , then $\pi \upharpoonright \rho$ denotes the suffix of π that remains after removing the prefix ρ , but not the last state of ρ .

The *length* of a path π , denoted $|\pi| \in \mathbb{N} \cup \{\infty\}$, is the number of transitions in π ; for instance $|s_0 s_1 s_2 s_3| = 3$.

A *completeness criterion* defines the *complete paths* as a subset of all paths. Complete paths model actual runs of the represented system.

Kripke structures

Let AP be a set of *atomic predicates*.

A *Kripke structure* over AP is tuple $(S, \rightarrow, \models)$ with

- S a set (of states),
- $\rightarrow \subseteq S \times S$, the *transition relation*, and
- $\models \subseteq S \times AP$.

$s \models p$ says that predicate $p \in AP$ *holds* in state $s \in S$.

A *path* in a Kripke structure is a nonempty finite or infinite sequence s_0, s_1, \dots of states that follows the transition relation.

If ρ is a finite prefix of π , then $\pi \upharpoonright \rho$ denotes the suffix of π that remains after removing the prefix ρ , but not the last state of ρ .

The *length* of a path π , denoted $|\pi| \in \mathbb{N} \cup \{\infty\}$, is the number of transitions in π ; for instance $|s_0 s_1 s_2 s_3| = 3$.

A *completeness criterion* defines the *complete paths* as a subset of all paths. Complete paths model actual runs of the represented system.

The default completeness criterion, called *progress*, classifies a path as complete iff it either is infinite or ends in a *deadlock state*.

LTL

Linear-time temporal logic (LTL) is designed to formulate useful properties of systems. Its syntax is

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{X}\varphi \mid \mathbf{F}\varphi \mid \mathbf{G}\varphi \mid \psi\mathbf{U}\varphi$$

with $p \in AP$ an atomic predicate. The propositional connectives \Rightarrow and \vee can be added as syntactic sugar. It is interpreted on paths.

LTL

Linear-time temporal logic (LTL) is designed to formulate useful properties of systems. Its syntax is

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{X}\varphi \mid \mathbf{F}\varphi \mid \mathbf{G}\varphi \mid \psi\mathbf{U}\varphi$$

with $p \in AP$ an atomic predicate. The propositional connectives \Rightarrow and \vee can be added as syntactic sugar. It is interpreted on paths.

$\pi \models \varphi$ says that the path π *satisfies* the formula φ , or that φ is *valid* on π .

LTL

Linear-time temporal logic (LTL) is designed to formulate useful properties of systems. Its syntax is

$$\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{X}\varphi \mid \mathbf{F}\varphi \mid \mathbf{G}\varphi \mid \psi\mathbf{U}\varphi$$

with $p \in AP$ an atomic predicate. The propositional connectives \Rightarrow and \vee can be added as syntactic sugar. It is interpreted on paths.

$\pi \models \varphi$ says that the path π *satisfies* the formula φ , or that φ is *valid* on π . This is inductively defined by

- $\pi \models p$, with $p \in AP$, iff $s \models p$, where s is the first state of π ,
- $\pi \models \neg\varphi$ iff $\pi \not\models \varphi$,
- $\pi \models \varphi \wedge \psi$ iff $\pi \models \varphi$ and $\pi \models \psi$,
- $\pi \models \mathbf{X}\varphi$ iff $|\pi| > 0$ and $\pi_{+1} \models \varphi$,
- $\pi \models \mathbf{F}\varphi$ iff $\pi \upharpoonright \rho \models \varphi$ for some finite prefix ρ of π ,
- $\pi \models \mathbf{G}\varphi$ iff $\pi \upharpoonright \rho \models \varphi$ for each finite prefix ρ of π , and
- $\pi \models \psi\mathbf{U}\varphi$ iff $\pi \upharpoonright \rho \models \varphi$ for some finite prefix ρ of π ,
and $\pi \upharpoonright \zeta \models \psi$ for each proper prefix ζ of ρ .

Interpreting LTL on states

Let CC be a completeness criteria.

For s a state in a Kripke structure, write

$P \models^{CC} \varphi$ iff $\pi \models \varphi$ for all complete paths starting in s .

Interpreting LTL on states

Let CC be a completeness criteria.

For s a state in a Kripke structure, write

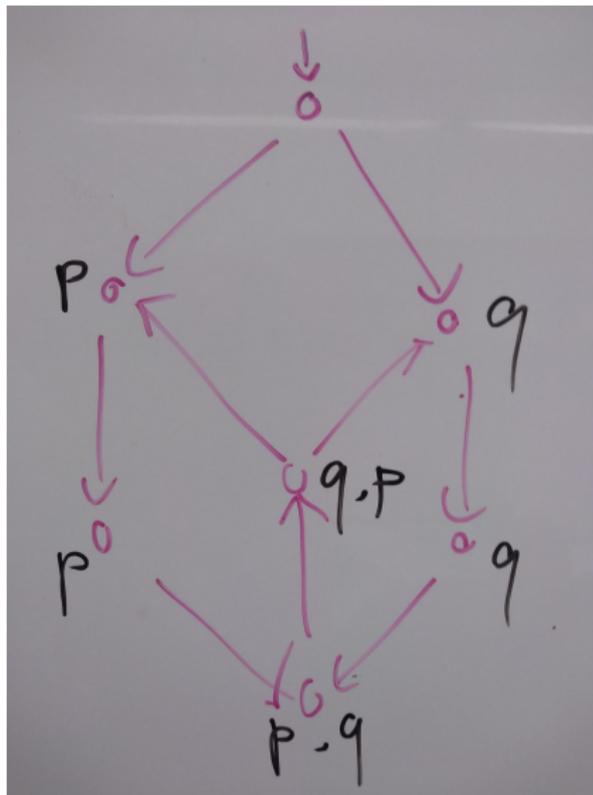
$P \models^{CC} \varphi$ iff $\pi \models \varphi$ for all complete paths starting in s .

We leave out the CC and write $P \models \varphi$ when taking *progress* as CC .

Homework

Which of these LTL formulae hold in the initial state of this Kripke structure?

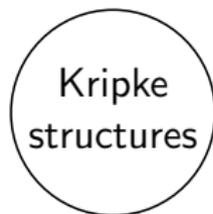
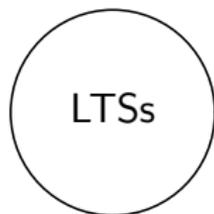
- (a) $\mathbf{G}(p \vee q)$
- (b) $\mathbf{FG}(p \vee q)$
- (c) $\mathbf{F}(\mathbf{G}p \vee \mathbf{G}q)$
- (d) $\mathbf{G}(q \rightarrow (q\mathbf{U}p))$
- (e) $\mathbf{G}((p \wedge \neg q) \rightarrow \mathbf{X}q)$
- (f) $\mathbf{G}((p \wedge q) \rightarrow \mathbf{X}(q\mathbf{U}p))$



De Nicola - Vaandrager translation

transitions labelled

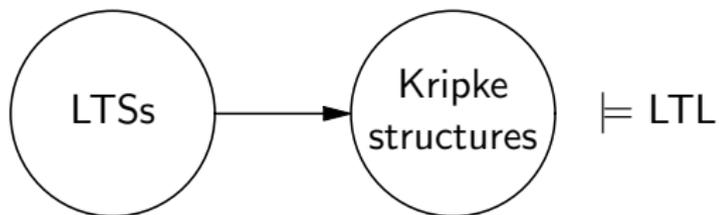
states labelled



\models LTL

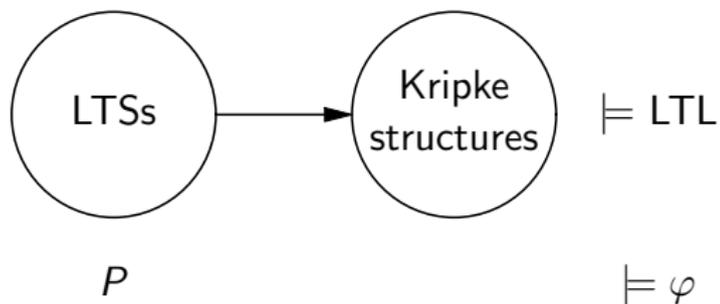
De Nicola - Vaandrager translation

transitions labelled states labelled



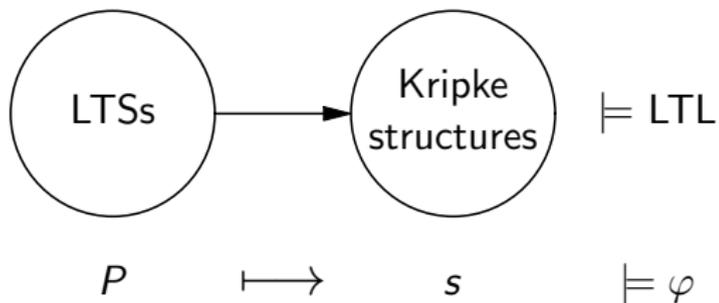
De Nicola - Vaandrager translation

transitions labelled states labelled



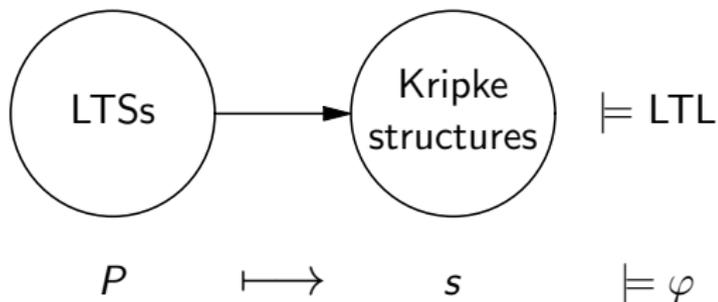
De Nicola - Vaandrager translation

transitions labelled states labelled



De Nicola - Vaandrager translation

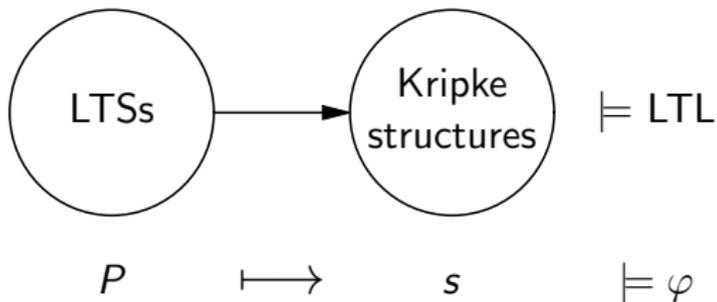
transitions labelled states labelled



Create a new state halfway along any transition labelled by a visible action, and move the transition label to that state.

De Nicola - Vaandrager translation

transitions labelled states labelled



Create a new state halfway along any transition labelled by a visible action, and move the transition label to that state.
 τ -transitions are not affected.

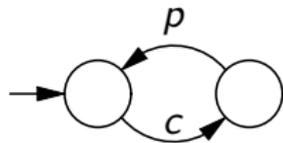
Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).

$$VM = c.p.VM$$

Pretzels

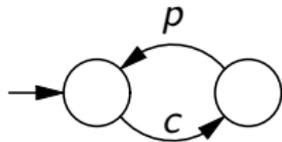
Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



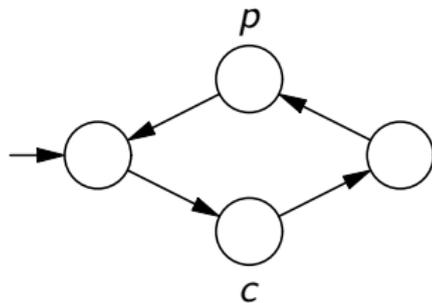
$$VM = c.p.VM$$

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).

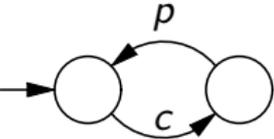


$$VM = c.p.VM$$

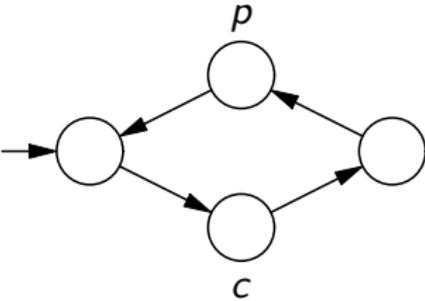


Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$

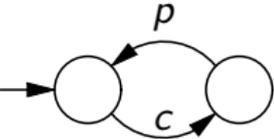


In LTL we have $VM \models G(c \Rightarrow Fp)$.

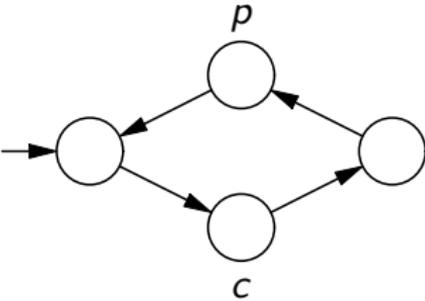
whenever a coin is inserted, eventually a pretzel is produced.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$



In LTL we have $VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$.

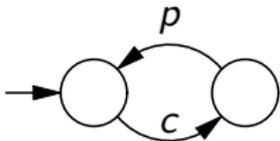
whenever a coin is inserted, eventually a pretzel is produced.

But we also have $VM \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$.

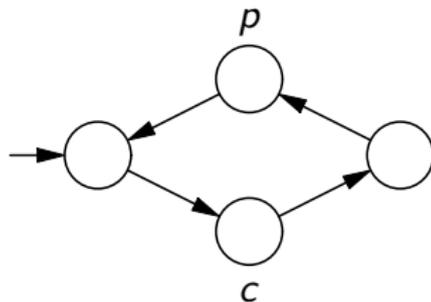
whenever a pretzel is produced, eventually a new coin will be inserted.

Pretzels

Consider a vending machine that alternately accepts a coin (c) and produces a pretzel (p).



$$VM = c.p.VM$$



In LTL we have $VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$.

whenever a coin is inserted, eventually a pretzel is produced.

But we also have $VM \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$.

whenever a pretzel is produced, eventually a new coin will be inserted.

This example shows that standard LTL is not suitable to correctly describe the behaviour of this vending machine.

Reactive Temporal Logic

Temporal judgements $P \models_B^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.

Reactive Temporal Logic

Temporal judgements $P \models_{B}^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.

CC is a *completeness criterion*, such as [progress](#), [justness](#), [weak fairness](#) or [strong fairness](#).

In default temporal logic infinite paths are complete; finite paths not.

B declares some of the finite paths as complete.

CC declares some of the infinite paths as incomplete.

Reactive Temporal Logic

Temporal judgements $P \models_{B}^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.

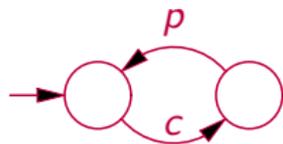
Reactive Temporal Logic

Temporal judgements $P \models_B^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.



$VM = c.p.VM$

$B = \{c\}$

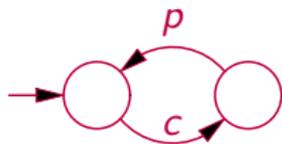
Reactive Temporal Logic

Temporal judgements $P \models_B^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.



$VM = c.p.VM$

$B = \{c\}$

In LTL we had $VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$.

whenever a coin is inserted, eventually a pretzel is produced.

But $VM \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$.

whenever a pretzel is produced, eventually a new coin will be inserted.

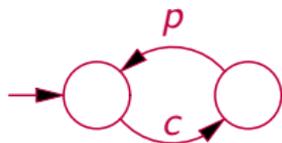
Reactive Temporal Logic

Temporal judgements $P \models_B^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.



$VM = c.p.VM$

$B = \{c\}$

In reactive LTL we have $VM \models_B \mathbf{G}(c \Rightarrow \mathbf{F}p)$.

whenever a coin is inserted, eventually a pretzel is produced.

But $VM \not\models_B \mathbf{G}(p \Rightarrow \mathbf{F}c)$.

whenever a pretzel is produced, eventually a new coin will be inserted.

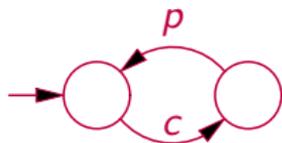
Reactive Temporal Logic

Temporal judgements $P \models_B^{CC} \varphi$

P is a process.

φ is a temporal formula in LTL, CTL, or some other logic.

B is a set of actions that may be blocked by the environment.



$VM = c.p.VM$

$B = \{c\}$

In reactive LTL we have $VM \models_B \mathbf{G}(c \Rightarrow \mathbf{F}p)$.

whenever a coin is inserted, eventually a pretzel is produced.

But $VM \not\models_B \mathbf{G}(p \Rightarrow \mathbf{F}c)$. We still have $VM \models_{\emptyset} \mathbf{G}(p \Rightarrow \mathbf{F}c)$.

whenever a pretzel is produced, eventually a new coin will be inserted.

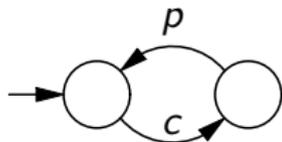
Reactive LTL

LTL formula are traditionally interpreted on infinite paths.

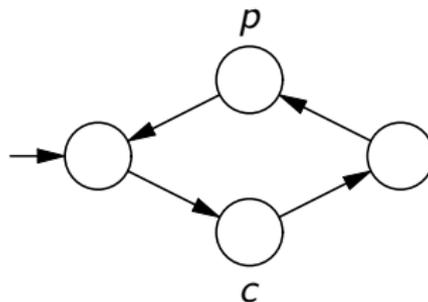
Reactive LTL

LTL formula are ~~traditionally~~ interpreted on ~~infinite~~ paths.

Reactive LTL

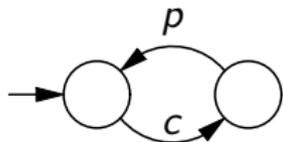


$$VM = c.p.VM$$

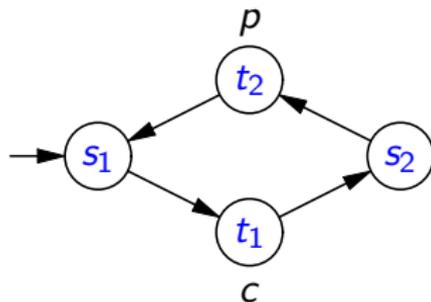


LTL formula are traditionally interpreted on infinite paths.

Reactive LTL

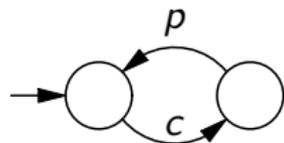


$$VM = c.p.VM$$

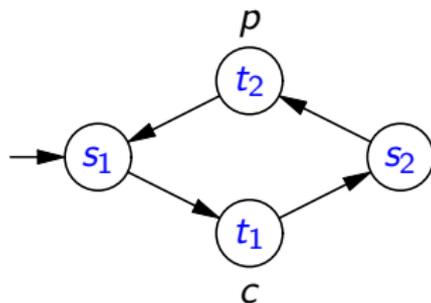


LTL formula are traditionally interpreted on infinite paths.

Reactive LTL



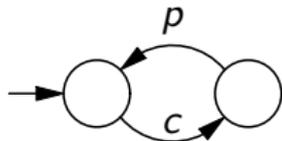
$$VM = c.p.VM$$



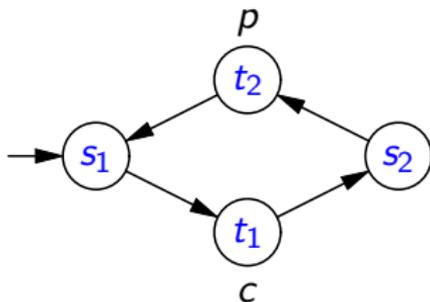
LTL formula are traditionally interpreted on infinite paths.

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \dots \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

Reactive LTL



$$VM = c.p.VM$$



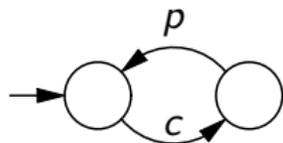
LTL formula are traditionally interpreted on infinite paths.

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \dots \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

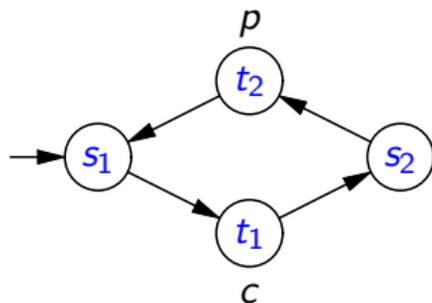
$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

$$VM \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

Reactive LTL



$$VM = c.p.VM$$



LTL formula are traditionally interpreted on infinite paths.

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \dots \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

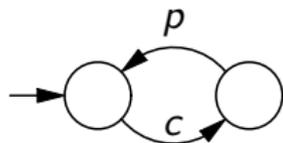
$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$$

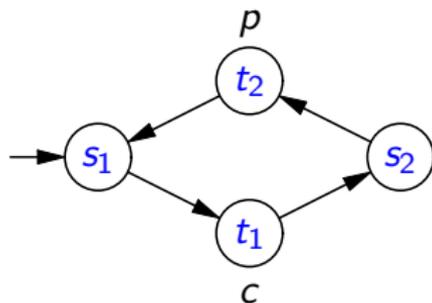
$$VM \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

$$VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$$

Reactive LTL



$$VM = c.p.VM$$



LTL formula are traditionally interpreted on infinite paths.

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \dots \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

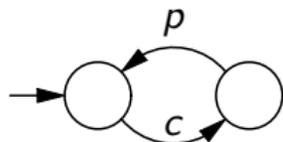
$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$$

$$s_1 t_1 s_2 t_2 s_1 t_1 s_2 \not\models \mathbf{G}(c \Rightarrow \mathbf{F}p)$$

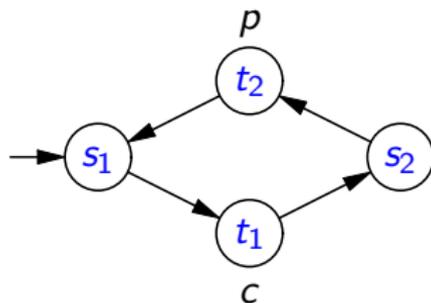
$$VM \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$$

$$VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$$

Reactive LTL



$$VM = c.p.VM$$



LTL formula are traditionally interpreted on infinite paths.

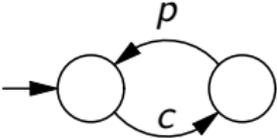
$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \dots \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$ complete

$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$ complete $VM \not\models_B \mathbf{G}(p \Rightarrow \mathbf{F}c)$

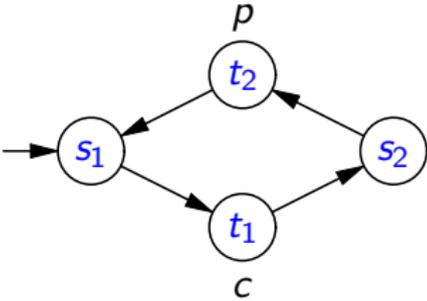
$s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$ complete $VM \models_B \mathbf{G}(c \Rightarrow \mathbf{F}p)$

$s_1 t_1 s_2 t_2 s_1 t_1 s_2 \not\models \mathbf{G}(c \Rightarrow \mathbf{F}p)$ incomplete

Reactive LTL



$$VM = c.p.VM$$



LTL formula are traditionally interpreted on infinite paths.

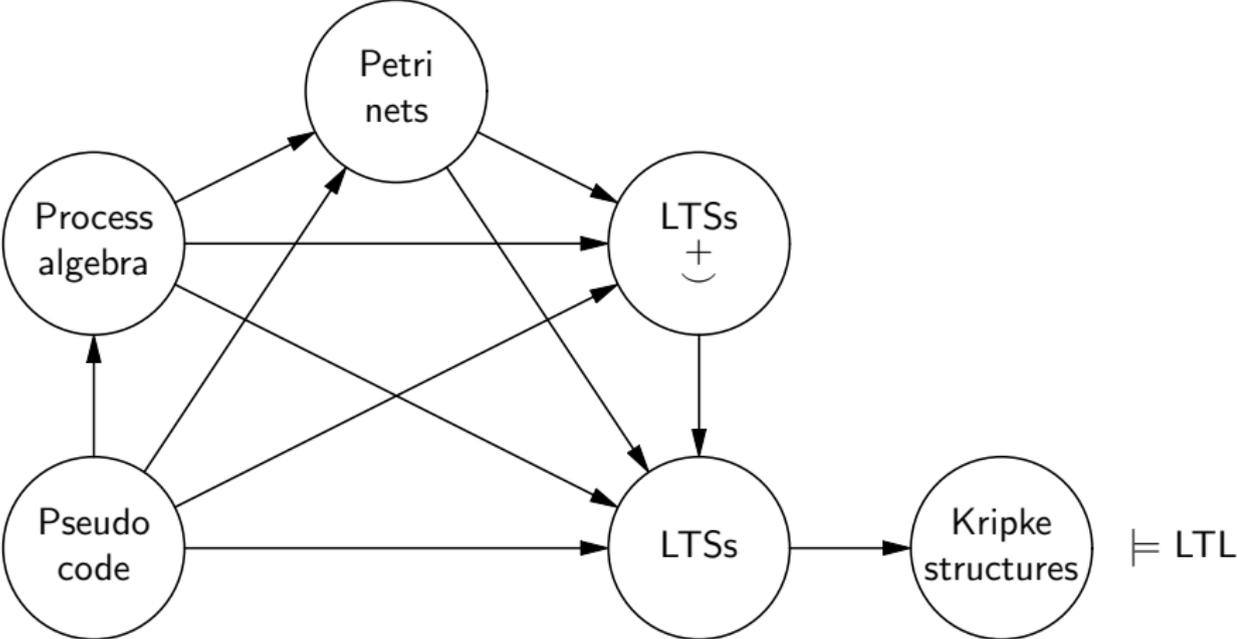
- $s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \dots \models \mathbf{G}(p \Rightarrow \mathbf{F}c)$ complete
- $s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$ complete $VM \not\models \mathbf{G}(p \Rightarrow \mathbf{F}c)$
- $s_1 t_1 s_2 t_2 s_1 t_1 s_2 t_2 s_1 \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$ complete $VM \models \mathbf{G}(c \Rightarrow \mathbf{F}p)$
- $s_1 t_1 s_2 t_2 s_1 t_1 s_2 \not\models \mathbf{G}(c \Rightarrow \mathbf{F}p)$ incomplete

$$B = \{c\}$$

$P \models \varphi$ iff $\pi \models \varphi$ for all complete paths π .

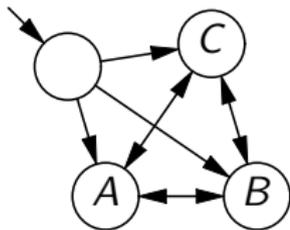
A path is **complete** iff either it is infinite, or its last state has outgoing transitions with labels from B only.

Models for distributed systems



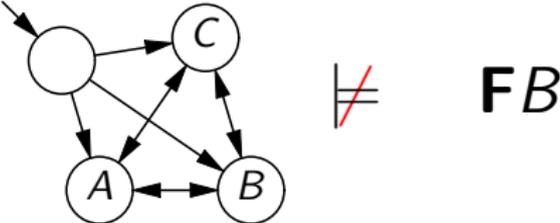
Fairness

Alice, Bart and Cameron stand behind a bar, ordering and drinking beer. As there is only one barman, they are served sequentially. None of them is served twice in a row, but each of them is ready for the next beer as soon as another person is served.



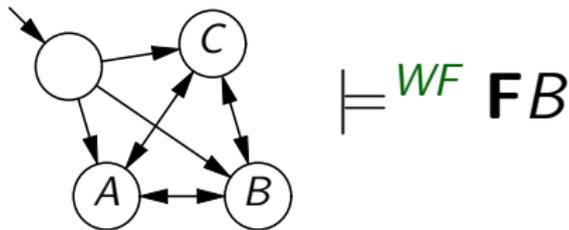
Fairness

Alice, Bart and Cameron stand behind a bar, ordering and drinking beer. As there is only one barman, they are served sequentially. None of them is served twice in a row, but each of them is ready for the next beer as soon as another person is served.



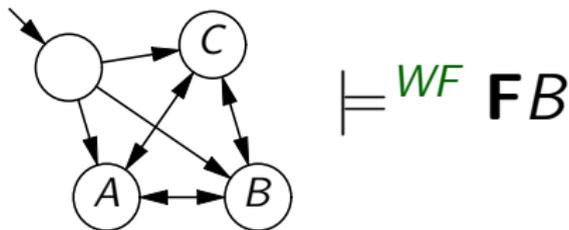
Fairness

Alice, Bart and Cameron stand behind a bar, ordering and drinking beer. As there is only one barman, they are served sequentially. None of them is served twice in a row, but each of them is ready for the next beer as soon as another person is served.



Fairness

Alice, Bart and Cameron stand behind a bar, ordering and drinking beer. As there is only one barman, they are served sequentially. None of them is served twice in a row, but each of them is ready for the next beer as soon as another person is served.



Assuming fairness often leads to unwarranted conclusions.

Progress

Bart is the only customer in a bar in London, with a single barman.
He only wants one beer.



Progress

Bart is the only customer in a bar in London, with a single barman.
He only wants one beer.



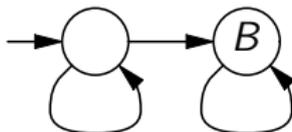
Justness

Bart is the only customer in a bar in London, with a single barman. He only wants one beer.

Alice and Cameron are in a bar in Tokyo. They drink a lot of beer.

There is no interaction of any kind between the two bars.

Yet, one may choose to model the drinking in these two bars as a single distributed system.



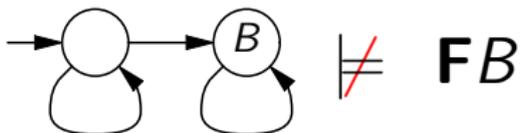
Justness

Bart is the only customer in a bar in London, with a single barman. He only wants one beer.

Alice and Cameron are in a bar in Tokyo. They drink a lot of beer.

There is no interaction of any kind between the two bars.

Yet, one may choose to model the drinking in these two bars as a single distributed system.



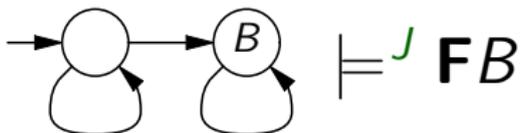
Justness

Bart is the only customer in a bar in London, with a single barman. He only wants one beer.

Alice and Cameron are in a bar in Tokyo. They drink a lot of beer.

There is no interaction of any kind between the two bars.

Yet, one may choose to model the drinking in these two bars as a single distributed system.



Each transition that is at some point enabled, will eventually occur, unless one of its necessary resources is used for another transition.

Session types

Another process algebra:

$$\begin{array}{l} T := \text{OK} \\ | \bigoplus_{i \in I} p_i! \lambda_i; T_i \\ | \sum_{i \in I} p_i? \lambda_i; T_i \\ | X \\ | \mu X. T \end{array} \qquad \begin{array}{l} N := p \llbracket T \rrbracket \\ | 0 \\ | N \parallel N \end{array}$$

Example:

$$\begin{array}{l} \text{seller1} \llbracket \mu X. \text{buyer1?}; X \rrbracket \\ \parallel \text{buyer1} \llbracket \mu Y. \text{seller1!}; Y \rrbracket \\ \parallel \text{seller2} \llbracket \mu Z. \text{buyer2?}; Z \rrbracket \\ \parallel \text{buyer2} \llbracket \mu W. \text{seller2!}; W \rrbracket \end{array}$$

Lock-freedom: **Everyone wishing to trade eventually does so.**

Session types

Another process algebra:

$$\begin{array}{l} T := \text{OK} \\ | \bigoplus_{i \in I} p_i ! \lambda_i ; T_i \\ | \sum_{i \in I} p_i ? \lambda_i ; T_i \\ | X \\ | \mu X . T \end{array} \qquad \begin{array}{l} N := p \llbracket T \rrbracket \\ | 0 \\ | N \parallel N \end{array}$$

Example:

$$\begin{array}{l} \text{seller1} \llbracket \mu X . \text{buyer1?}; X \rrbracket \\ \parallel \text{buyer1} \llbracket \mu Y . \text{seller1!}; Y \rrbracket \\ \parallel \text{seller2} \llbracket \mu Z . \text{buyer2?}; Z \rrbracket \\ \parallel \text{buyer2} \llbracket \mu W . \text{seller2!}; W \rrbracket \end{array}$$

Lock-freedom: **Everyone wishing to trade eventually does so.**

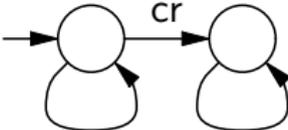
Theorem: $\mathcal{L}(J)$ holds iff network is well-typed and race-free.

Justness

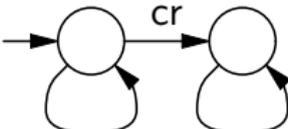
Cateline, eating a croissant in Paris:



In parallel with Alice, making an infinite sequence of phone calls in London:



Cateline, making an infinite sequence of choices between eating her croissant and making yet another phone call:



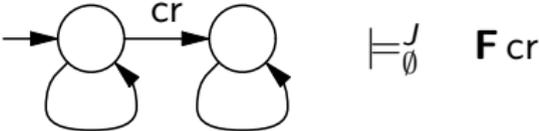
In Petri nets, these systems have totally different representations.

Justness

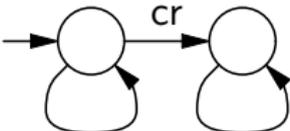
Cateline, eating a croissant in Paris:



In parallel with Alice, making an infinite sequence of phone calls in London:



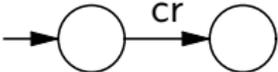
Cateline, making an infinite sequence of choices between eating her croissant and making yet another phone call:



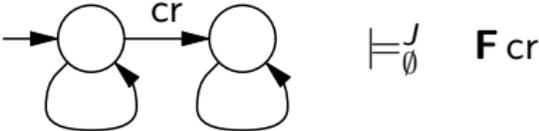
In Petri nets, these systems have totally different representations.

Justness

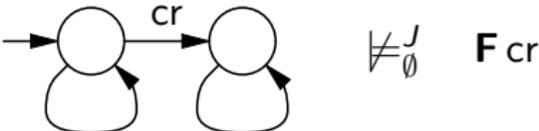
Cateline, eating a croissant in Paris:



In parallel with Alice, making an infinite sequence of phone calls in London:

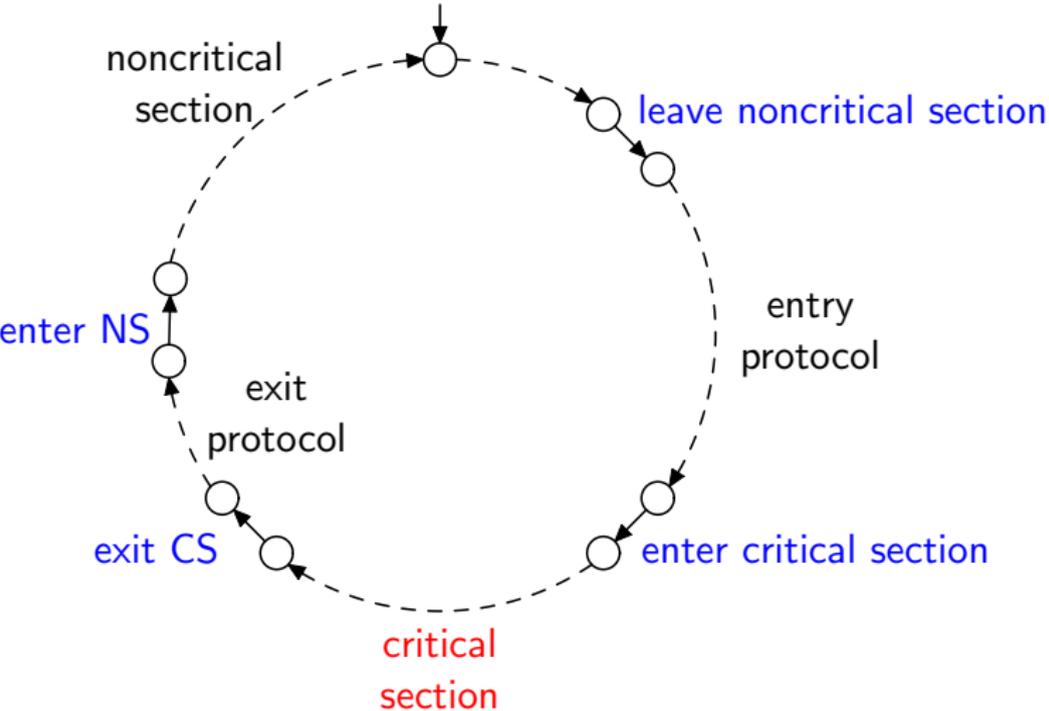


Cateline, making an infinite sequence of choices between eating her croissant and making yet another phone call:



In Petri nets, these systems have totally different representations.

Mutual exclusion



Peterson's mutual exclusion algorithm

Process A

repeat forever

$\left\{ \begin{array}{l} \ell_1 \text{ leave noncritical section} \\ \ell_2 \text{ readyA} := \text{true} \\ \ell_3 \text{ turn} := B \\ \ell_4 \text{ await} (\text{readyB} = \text{false} \vee \text{turn} = A) \\ \ell_5 \text{ enter critical section} \\ \ell_6 \text{ leave critical section} \\ \ell_7 \text{ readyA} := \text{false} \\ \ell_8 \text{ enter noncritical section} \end{array} \right.$

Process B

repeat forever

$\left\{ \begin{array}{l} m_1 \text{ leave noncritical section} \\ m_2 \text{ readyB} := \text{true} \\ m_3 \text{ turn} := A \\ m_4 \text{ await} (\text{readyA} = \text{false} \vee \text{turn} = B) \\ m_5 \text{ enter critical section} \\ m_6 \text{ leave critical section} \\ m_7 \text{ readyB} := \text{false} \\ m_8 \text{ enter noncritical section} \end{array} \right.$

Peterson's algorithm (pseudocode)

Peterson's mutual exclusion algorithm

Process A

repeat forever

$\left\{ \begin{array}{l} \ell_1 \text{ leave noncritical section} \\ \ell_2 \text{ readyA} := \text{true} \\ \ell_3 \text{ turn} := B \\ \ell_4 \text{ await} (\text{readyB} = \text{false} \vee \text{turn} = A) \\ \ell_5 \text{ enter critical section} \\ \ell_6 \text{ leave critical section} \\ \ell_7 \text{ readyA} := \text{false} \\ \ell_8 \text{ enter noncritical section} \end{array} \right.$

Process B

repeat forever

$\left\{ \begin{array}{l} m_1 \text{ leave noncritical section} \\ m_2 \text{ readyB} := \text{true} \\ m_3 \text{ turn} := A \\ m_4 \text{ await} (\text{readyA} = \text{false} \vee \text{turn} = B) \\ m_5 \text{ enter critical section} \\ m_6 \text{ leave critical section} \\ m_7 \text{ readyB} := \text{false} \\ m_8 \text{ enter noncritical section} \end{array} \right.$

Peterson's algorithm (pseudocode)

Mutual exclusion: Never are both processes together in the critical section.

Starvation-freedom: Any process that leaves its noncritical section, will enter its critical section.

Peterson's mutual exclusion algorithm

Process A

repeat forever

$\left\{ \begin{array}{l} \ell_1 \text{ leave noncritical section} \\ \ell_2 \text{ readyA} := \text{true} \\ \ell_3 \text{ turn} := B \\ \ell_4 \text{ await} (\text{readyB} = \text{false} \vee \text{turn} = A) \\ \ell_5 \text{ enter critical section} \\ \ell_6 \text{ leave critical section} \\ \ell_7 \text{ readyA} := \text{false} \\ \ell_8 \text{ enter noncritical section} \end{array} \right.$

Process B

repeat forever

$\left\{ \begin{array}{l} m_1 \text{ leave noncritical section} \\ m_2 \text{ readyB} := \text{true} \\ m_3 \text{ turn} := A \\ m_4 \text{ await} (\text{readyA} = \text{false} \vee \text{turn} = B) \\ m_5 \text{ enter critical section} \\ m_6 \text{ leave critical section} \\ m_7 \text{ readyB} := \text{false} \\ m_8 \text{ enter noncritical section} \end{array} \right.$

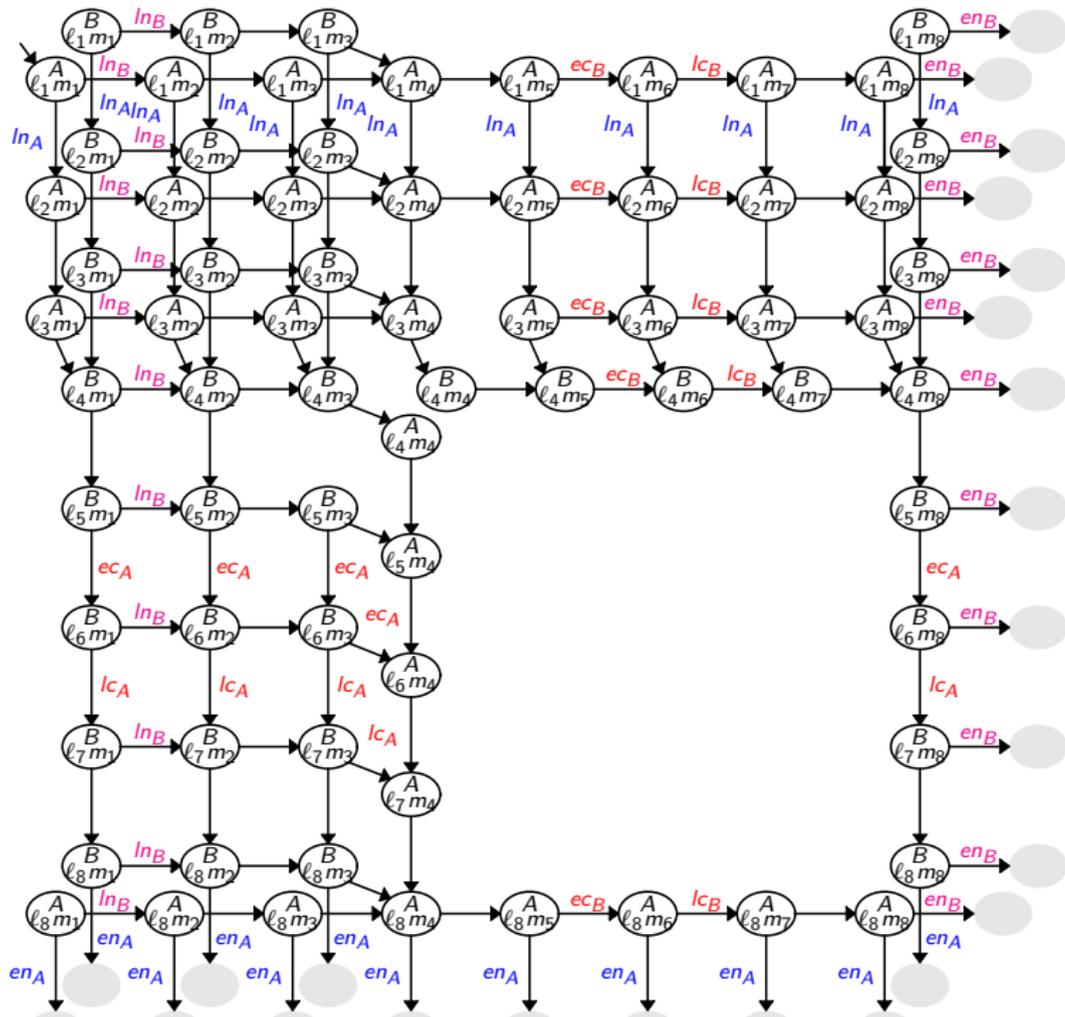
Peterson's algorithm (pseudocode)

Mutual exclusion: Never are both processes together in the critical section.

Starvation-freedom: Any process that leaves its noncritical section, will enter its critical section.

Speed independence: the protocol should work regardless how long things take.

Atomicity: a memory cell cannot handle a read and a write at the same time. So a write will have to wait until a pending read is finished.



Correctness properties of MUTEX protocols

$$\begin{aligned} \text{(ORD)} \quad P \models & ((\neg act_i) \mathbf{W} ln_i) \wedge \mathbf{G} (ln_i \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} ec_i)) \\ & \wedge \mathbf{G} (ec_i \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} lc_i)) \\ & \wedge \mathbf{G} (lc_i \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} en_i)) \\ & \wedge \mathbf{G} (en_i \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} ln_i)) \end{aligned}$$

for $i = 1, \dots, N$. Here $act_i := (ln_i \vee ec_i \vee lc_i \vee en_i)$.

This requirement says that in a mutex protocol the actions ln_i , ec_i , lc_i and en_i must occur in the right order.

Correctness properties of MUTEX protocols

$$\begin{aligned} \text{(ORD)} \quad P \models & ((\neg act_i) \mathbf{W} ln_i) \wedge \mathbf{G} (ln_i \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} ec_j)) \\ & \wedge \mathbf{G} (ec_j \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} lc_j)) \\ & \wedge \mathbf{G} (lc_j \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} en_j)) \\ & \wedge \mathbf{G} (en_j \Rightarrow \mathbf{Y}((\neg act_i) \mathbf{W} ln_i)) \end{aligned}$$

for $i = 1, \dots, N$. Here $act_i := (ln_i \vee ec_j \vee lc_j \vee en_j)$.

$$\text{(ME)} \quad P \models \mathbf{G} (ec_i \Rightarrow ((\neg ec_j) \mathbf{W} lc_i))$$

for all $i, j = 1, \dots, N$ with $i \neq j$.

Correctness properties of MUTEX protocols

$$(EC^J) \quad P \models_B^J \mathbf{G}(ln_i \Rightarrow \mathbf{F}ec_i)$$

for $i = 1, \dots, N$.

Correctness properties of MUTEX protocols

$$(EC^J) \quad P \models_B^J \mathbf{G}(ln_i \Rightarrow \mathbf{F}ec_i)$$

for $i = 1, \dots, N$.

$$(LC^J) \quad P \models_B^J \mathbf{G}(ec_i \Rightarrow \mathbf{F}lc_i)$$

$$(EN^J) \quad P \models_B^J \mathbf{G}(lc_i \Rightarrow \mathbf{F}en_i)$$

$$(LN^J) \quad P \models_{B \setminus \{ln_i\}}^J \mathbf{F}ln_i \wedge \mathbf{G}(en_i \Rightarrow \mathbf{F}ln_i)$$

for $i = 1, \dots, N$.

Memory model

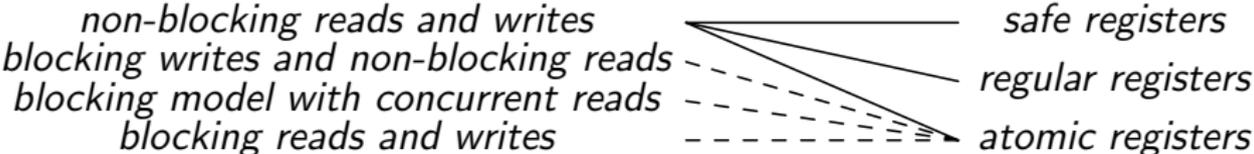
The processes in a mutex protocol communicate with each other by writing to and reading from shared registers.

Memory model

The processes in a mutex protocol communicate with each other by writing to and reading from shared registers. Whether a mutex protocol is correct depends on the way these registers work.

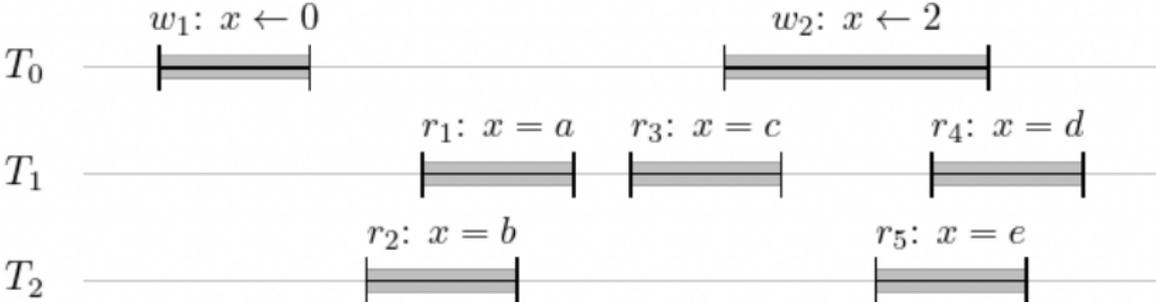
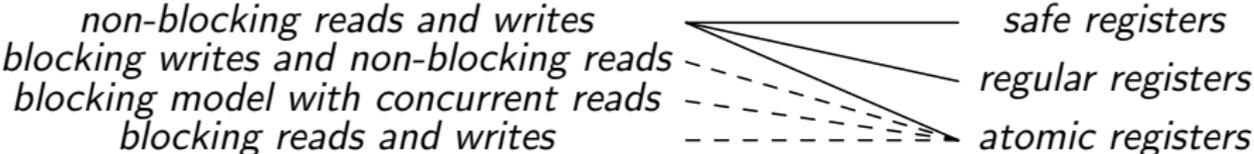
Memory model

The processes in a mutex protocol communicate with each other by writing to and reading from shared registers. Whether a mutex protocol is correct depends on the way these registers work.



Memory model

The processes in a mutex protocol communicate with each other by writing to and reading from shared registers. Whether a mutex protocol is correct depends on the way these registers work.



Model checking

Traditional verification of mutual exclusion algorithms:

pen-and-paper proofs using behavioural reasoning.

“the behavioral reasoning used in our correctness proofs, and in most other published correctness proofs of concurrent algorithms, is inherently unreliable” [Lamport 1986]

This is especially the case when dealing with non-atomic registers.

Model checking

Traditional verification of mutual exclusion algorithms:

pen-and-paper proofs using behavioural reasoning.

“the behavioral reasoning used in our correctness proofs, and in most other published correctness proofs of concurrent algorithms, is inherently unreliable” [Lamport 1986]

This is especially the case when dealing with non-atomic registers.

Alternatives: interactive theorem proving or model checking.

Model checking: Precise modelling requires great care, the verification requires a mere button-push and some patience.

Model checking

Traditional verification of mutual exclusion algorithms:

pen-and-paper proofs using behavioural reasoning.

“the behavioral reasoning used in our correctness proofs, and in most other published correctness proofs of concurrent algorithms, is inherently unreliable” [Lamport 1986]

This is especially the case when dealing with non-atomic registers.

Alternatives: interactive theorem proving or model checking.

Model checking: Precise modelling requires great care, the verification requires a mere button-push and some patience.

But only finite state spaces.

So no bakery protocol

and we can check only for a small number of threads.

Results

<i>Algorithm</i>	<i># threads</i>	<i>Safe</i>	<i>Regular</i>	<i>Atomic</i>			
		<i>T</i>	<i>T</i>	<i>T</i>	<i>S</i>	<i>I</i>	<i>A</i>
Anderson	2	S	S	S	S	M	M
Aravind BLRU	3	S	S	S	M	M	M
Aravind BLRU (alt.)	3	S	S	S	S	M	M
Attiya-Welch (orig.)	2	D	S	S	D	M	M
Attiya-Welch (orig., alt.)	2	S	S	S	D	M	M
Attiya-Welch (var.)	2	M	M	S	D	M	M
Attiya-Welch (var., alt.)	2	S	S	S	D	M	M
Burns-Lynch	3	D	D	D	D	M	M
Dekker	2	M	M	S	D	M	M
Dekker (alt.)	2	M	M	S	S	M	M
Dekker RW-safe	2	S	S	S	D	M	M
Dekker RW-safe (DFtoSF)	2	S	S	S	S	M	M
Dijkstra	3	M	D	D	M	M	M

Results (continued)

Kessels	2	X	X	S	S	M	M
Knuth	3	M	S	S	M	M	M
Lamport 1-bit	3	D	D	D	D	M	M
Lamport 1-bit (DFtoSF)	3	S	S	S	S	M	M
Lamport 3-bit	3	S	S	S	S	M	M
Peterson	2	X	X	S	S	M	M
Szymanski flag (int.)	3	X	X	S	S	M	M
Szymanski flag (bit)	3	X	X	X	X	X	X
Szymanski 3-bit lin. wait	3	X	X	X	X	X	X
Szymanski 3-bit lin. wait (alt.)	2	S	S	S	S	M	M

Speed independence

Nothing may be assumed about the relative speed of the processes competing for access to the critical section. [Dijkstra'65]

Speed independence

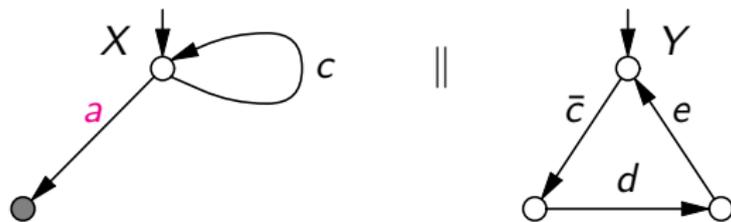
Nothing may be assumed about the relative speed of the processes competing for access to the critical section. [Dijkstra'65]

If two processes A and B are engaged in a race, and A has nothing else to do but performing the winning action, whilst B has a long list of tasks that must be done first, it may still happen that B wins.

Speed independence

Nothing may be assumed about the relative speed of the processes competing for access to the critical section. [Dijkstra'65]

CCS process $(X|Y)\backslash c$ with $X \stackrel{\text{def}}{=} a.\mathbf{0} + c.X$ and $Y \stackrel{\text{def}}{=} \bar{c}.d.e.Y$.

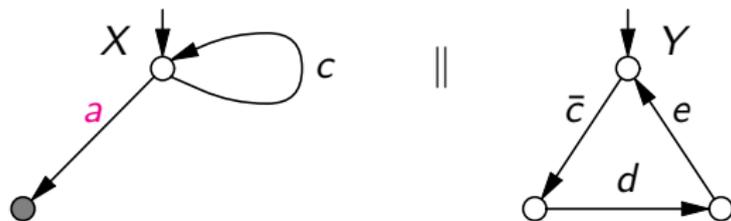


If two processes A and B are engaged in a race, and A has nothing else to do but performing the winning action, whilst B has a long list of tasks that must be done first, it may still happen that B wins.

Speed independence

Nothing may be assumed about the relative speed of the processes competing for access to the critical section. [Dijkstra'65]

CCS process $(X|Y)\backslash c$ with $X \stackrel{\text{def}}{=} a.\mathbf{0} + c.X$ and $Y \stackrel{\text{def}}{=} \bar{c}.d.e.Y$.



Here it is possible that a never happens.

If two processes A and B are engaged in a race, and A has nothing else to do but performing the winning action, whilst B has a long list of tasks that must be done first, it may still happen that B wins.

Impossibility of speed ind. mutual excl. under atomicity

For mutual exclusion to be possible, there must be a variable *ready P* that is written by Proc. P to request entry to CS.

Impossibility of speed ind. mutual excl. under atomicity

For mutual exclusion to be possible, there must be a variable *readyP* that is written by Proc. *P* to request entry to CS.

readyP must be read by Proc. *Q*, before Proc. *Q* can enter CS.

Impossibility of speed ind. mutual excl. under atomicity

For mutual exclusion to be possible, there must be a variable *readyP* that is written by Proc. *P* to request entry to CS.

readyP must be read by Proc. *Q*, before Proc. *Q* can enter CS.

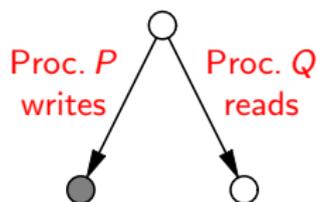


It suffices to present a scenario where Proc. *P* is ready to write to *readyP* yet never succeeds in doing so, as that would violate **EC** or **LN**.

Impossibility of speed ind. mutual excl. under atomicity

For mutual exclusion to be possible, there must be a variable *readyP* that is written by Proc. *P* to request entry to CS.

readyP must be read by Proc. *Q*, before Proc. *Q* can enter CS.

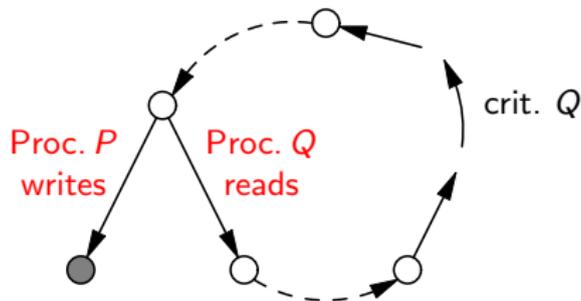


It suffices to present a scenario where Proc. *P* is ready to write to *readyP* yet never succeeds in doing so, as that would violate **EC** or **LN**.

Impossibility of speed ind. mutual excl. under atomicity

For mutual exclusion to be possible, there must be a variable *readyP* that is written by Proc. *P* to request entry to CS.

readyP must be read by Proc. *Q*, before Proc. *Q* can enter CS.



It suffices to present a scenario where Proc. *P* is ready to write to *readyP* yet never succeeds in doing so, as that would violate **EC** or **LN**.