Overview

- Introduction
- Modelling parallel systems
- **Linear Time Properties**
  - state-based and linear time view
  - definition of linear time properties
  - invariants and safety
  - liveness and fairness ⟵
- Regular Properties
- Linear Temporal Logic
- Computation-Tree Logic
- Equivalences and Abstraction

Liveness

“liveness: something good will happen.”

"liveness: something good will happen."

"event a will occur eventually"

Liveness

"liveness: something good will happen."

"event *a* will occur eventually"

e.g., termination for sequential programs

"liveness: something good will happen."

"event ***a*** will occur eventually"

e.g., termination for sequential programs

"event ***a*** will occur infinitely many times"

e.g., starvation freedom for dining philosophers

Liveness

"liveness: something good will happen."

"event ***a*** will occur eventually"

e.g., termination for sequential programs

"event ***a*** will occur infinitely many times"

e.g., starvation freedom for dining philosophers

"whenever event ***b*** occurs then event ***a***
will occur sometimes in the future"

“liveness: something good will happen.”

“event ***a*** will occur eventually”

e.g., termination for sequential programs

“event ***a*** will occur infinitely many times”

e.g., starvation freedom for dining philosophers

“whenever event ***b*** occurs then event ***a*** will occur sometimes in the future”

e.g., every waiting process enters eventually its critical section

which property type?

- Each philosopher thinks infinitely often.

which property type?

- Each philosopher thinks infinitely often.

liveness

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

  invariant

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

  invariant

- Whenever a philosopher eats then he has been thinking at some time before.

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

  invariant

- Whenever a philosopher eats then he has been thinking at some time before.

  safety

which property type?

- Each philosopher thinks infinitely often.

liveness

- Two philosophers next to each other never eat at the same time.

invariant

- Whenever a philosopher eats then he has been thinking at some time before.

safety

- Whenever a philosopher eats then he will think some time afterwards.

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

  invariant

- Whenever a philosopher eats then he has been thinking at some time before.

  safety

- Whenever a philosopher eats then he will think some time afterwards.

  liveness

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

  invariant

- Whenever a philosopher eats then he has been thinking at some time before.

  safety

- Whenever a philosopher eats then he will think some time afterwards.

  liveness

- Between two eating phases of philosopher i lies at least one eating phase of philosopher $i+1$.

which property type?

- Each philosopher thinks infinitely often.

  liveness

- Two philosophers next to each other never eat at the same time.

  invariant

- Whenever a philosopher eats then he has been thinking at some time before.

  safety

- Whenever a philosopher eats then he will think some time afterwards.

  liveness

- Between two eating phases of philosopher i lies at least one eating phase of philosopher $i+1$.

  safety

Liveness

many different formal definitions of liveness
have been suggested in the literature

Liveness

many different formal definitions of liveness
have been suggested in the literature

here: one just example for a formal definition
of liveness

Definition of liveness properties

LF2.6-DEF-LIVENESS

Definition of liveness properties

Let E be an LT property over AP, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$.

E is called a liveness property if each finite word over AP can be extended to an infinite word in E

Definition of liveness properties

Let E be an LT property over AP, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$.

E is called a liveness property if each finite word over AP can be extended to an infinite word in E, i.e., if

$$pref(E) = \left(2^{AP}\right)^{+}$$

recall: $pref(E) =$ set of all finite, nonempty prefixes of words in E

Definition of liveness properties

Let E be an LT property over AP, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$.

E is called a liveness property if each finite word over AP can be extended to an infinite word in E, i.e., if

$$pref(E) \; = \; \left(2^{AP}\right)^{+}$$

Examples:

- each process will eventually enter its critical section
- each process will enter its critical section infinitely often
- whenever a process has requested its critical section then it will eventually enter its critical section

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = \left(2^{AP}\right)^{+}$

Examples for $AP = \{crit_i : i = 1, \ldots, n\}$:

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = (2^{AP})^+$

Examples for $AP = \{crit_i : i = 1, \ldots, n\}$:

- each process will eventually enter its critical section

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = (2^{AP})^+$

Examples for $AP = \{crit_i : i = 1, \ldots, n\}$:

- each process will eventually enter its critical section

E = set of all infinite words $A_0 A_1 A_2 \ldots$ s.t.

$$\forall i \in \{1, \ldots, n\} \; \exists k \geq 0. \; crit_i \in A_k$$

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = \left(2^{AP}\right)^+$

Examples for $AP = \{crit_i : i = 1, \ldots, n\}$:

- each process will eventually enter its critical section
- each process will enter its critical section infinitely often

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = (2^{AP})^+$

Examples for $AP = \{crit_i : i = 1, \ldots, n\}$:

- each process will eventually enter its critical section
- each process will enter its critical section infinitely often

E = set of all infinite words $A_0 A_1 A_2 \ldots$ s.t.

$$\forall i \in \{1, \ldots, n\} \ \overset{\infty}{\exists} k \geq 0.\ crit_i \in A_k$$

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = (2^{AP})^+$

Examples for $AP = \{wait_i, crit_i : i = 1, \ldots, n\}$:

- each process will eventually enter its critical section
- each process will enter its crit. section inf. often
- whenever a process is waiting then it will eventually enter its critical section

Examples for liveness properties

An LT property E over AP is called a liveness property if $pref(E) = (2^{AP})^+$

Examples for $AP = \{wait_i, crit_i : i = 1, \ldots, n\}$:

- each process will eventually enter its critical section
- each process will enter its crit. section inf. often
- whenever a process is waiting then it will eventually enter its critical section

E = set of all infinite words $A_0\, A_1\, A_2 \ldots$ s.t.

$$\forall i \in \{1, \ldots, n\}\ \forall j \geq 0.\ wait_i \in A_j \longrightarrow \exists k > j.\ crit_i \in A_k$$

Recall: safety properties, prefix closure

Let E be an LT-property, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$

Recall: safety properties, prefix closure

Let E be an LT-property, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$

E is a safety property

iff $\forall \sigma \in \left(2^{AP}\right)^{\omega} \setminus E \;\; \exists A_0 A_1 \ldots A_n \in pref(\sigma)$ s.t.

$$\left\{\sigma' \in E : A_0 A_1 \ldots A_n \in pref(\sigma')\right\} = \varnothing$$

Recall: safety properties, prefix closure

Let E be an LT-property, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$

E is a safety property

iff $\forall \sigma \in \left(2^{AP}\right)^{\omega} \setminus E \;\; \exists A_0 A_1 \ldots A_n \in pref(\sigma)$ s.t.

$$\left\{\sigma' \in E : A_0 A_1 \ldots A_n \in pref(\sigma')\right\} = \varnothing$$

remind:

$pref(\sigma)$ = set of all finite, nonempty prefixes of σ

$$pref(E) = \bigcup_{\sigma \in E} pref(\sigma)$$

Recall: safety properties, prefix closure

Let E be an LT-property, i.e., $E \subseteq \left(2^{AP}\right)^{\omega}$

E is a safety property

iff $\forall \sigma \in \left(2^{AP}\right)^{\omega} \setminus E \;\; \exists A_0 A_1 \ldots A_n \in pref(\sigma)$ s.t.

$$\left\{\sigma' \in E : A_0 A_1 \ldots A_n \in pref(\sigma')\right\} = \varnothing$$

iff $cl(E) = E$

remind: $cl(E) = \left\{\sigma \in \left(2^{AP}\right)^{\omega} : pref(\sigma) \subseteq pref(E)\right\}$

$pref(\sigma)$ = set of all finite, nonempty prefixes of σ

$$pref(E) = \bigcup_{\sigma \in E} pref(\sigma)$$

Decomposition theorem

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof:

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

---

remind: $cl(E) = \{\sigma \in (2^{AP})^\omega : pref(\sigma) \subseteq pref(E)\}$

$pref(\sigma)$ = set of all finite, nonempty prefixes of σ

$$pref(E) = \bigcup_{\sigma \in E} pref(\sigma)$$

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

$$LIVE \stackrel{\text{def}}{=} E \cup \left((2^{AP})^{\omega} \setminus cl(E) \right)$$

---

remind: $cl(E) = \left\{ \sigma \in (2^{AP})^{\omega} : pref(\sigma) \subseteq pref(E) \right\}$

$pref(\sigma)$ = set of all finite, nonempty prefixes of σ

$$pref(E) = \bigcup_{\sigma \in E} pref(\sigma)$$

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

$$LIVE \stackrel{\text{def}}{=} E \cup \big((2^{AP})^\omega \setminus cl(E) \big)$$

Show that:

- $E = SAFE \cap LIVE$
- $SAFE$ is a safety property
- $LIVE$ is a liveness property

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

$$LIVE \stackrel{\text{def}}{=} E \cup \left((2^{AP})^\omega \setminus cl(E) \right)$$

Show that:

- $E = SAFE \cap LIVE$ $\checkmark$
- $SAFE$ is a safety property
- $LIVE$ is a liveness property

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

$$LIVE \stackrel{\text{def}}{=} E \cup \left(\left(2^{AP}\right)^{\omega} \setminus cl(E) \right)$$

Show that:

- $E = SAFE \cap LIVE$ $\surd$
- $SAFE$ is a safety property as $cl(SAFE) = SAFE$
- $LIVE$ is a liveness property

Decomposition theorem

For each LT-property E, there exists a safety property $SAFE$ and a liveness property $LIVE$ s.t.

$$E = SAFE \cap LIVE$$

Proof: Let $SAFE \stackrel{\text{def}}{=} cl(E)$

$$LIVE \stackrel{\text{def}}{=} E \cup \left(\left(2^{AP}\right)^{\omega} \setminus cl(E) \right)$$

Show that:

- $E = SAFE \cap LIVE$ $\surd$
- $SAFE$ is a safety property as $cl(SAFE) = SAFE$
- $LIVE$ is a liveness property, i.e., $pref(LIVE) = \left(2^{AP}\right)^{+}$

Being safe and live

Which LT properties are both
a safety and a liveness property?

Being safe and live

Which LT properties are both
a safety and a liveness property?

answer: The set $(2^{AP})^{\omega}$ is the only LT property which is a safety property and a liveness property

Which LT properties are both a safety and a liveness property?

answer: The set $(2^{AP})^{\omega}$ is the only LT property which is a safety property and a liveness property

- $(2^{AP})^{\omega}$ is a safety and a liveness property: $\surd$

> Which LT properties are both
> a safety and a liveness property?

answer: The set $\left(2^{AP}\right)^{\omega}$ is the only LT property which is a safety property and a liveness property

- $\left(2^{AP}\right)^{\omega}$ is a safety and a liveness property: √
- If E is a liveness property then

$$pref(E) = \left(2^{AP}\right)^{+}$$

Being safe and live

Which LT properties are both
a safety and a liveness property?

answer: The set $(2^{AP})^\omega$ is the only LT property which is a safety property and a liveness property

- $(2^{AP})^\omega$ is a safety and a liveness property: $\surd$
- If E is a liveness property then

$$pref(E) = (2^{AP})^+$$
$$\Longrightarrow \quad cl(E) = (2^{AP})^\omega$$

Being safe and live

> Which LT properties are both
> a safety and a liveness property?

answer: The set $(2^{AP})^\omega$ is the only LT property which is a safety property and a liveness property

- $(2^{AP})^\omega$ is a safety and a liveness property: √
- If E is a liveness property then

$$pref(E) = (2^{AP})^+$$

$$\Longrightarrow \quad cl(E) = (2^{AP})^\omega$$

If E is a safety property too, then $cl(E) = E$.

Being safe and live

Which LT properties are both
a safety and a liveness property?

answer: The set $(2^{AP})^{\omega}$ is the only LT property which is a safety property and a liveness property

- $(2^{AP})^{\omega}$ is a safety and a liveness property: √
- If E is a liveness property then

$$\begin{aligned} & pref(E) = (2^{AP})^{+} \\ \Longrightarrow \quad & cl(E) = (2^{AP})^{\omega} \end{aligned}$$

If E is a safety property too, then $cl(E) = E$.
Hence $E = cl(E) = (2^{AP})^{\omega}$.

Observation

liveness properties are often violated
although we expect them to hold

Two independent traffic lights

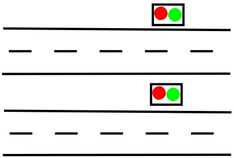

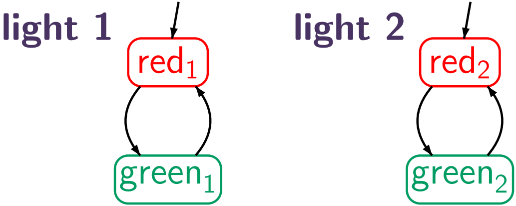

Two independent traffic lights

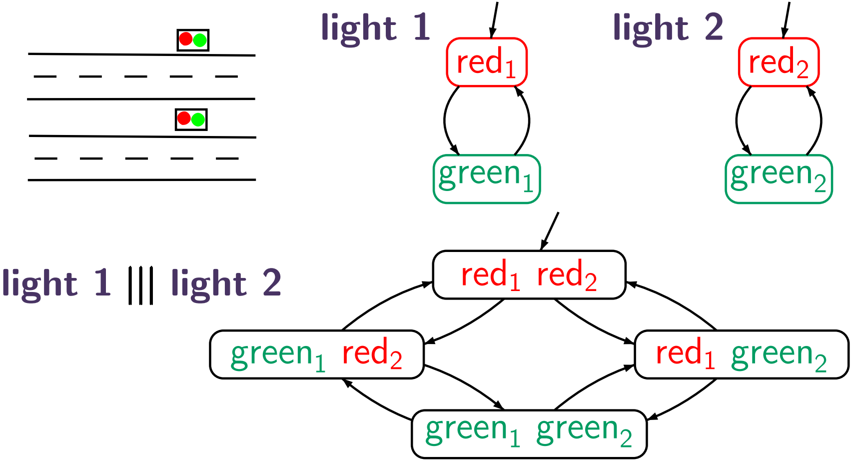

Two independent traffic lights

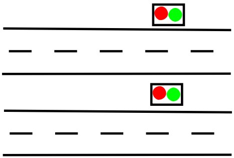

light 1

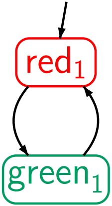

light 2

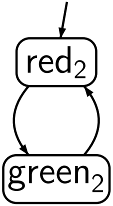

light 1 ||| light 2

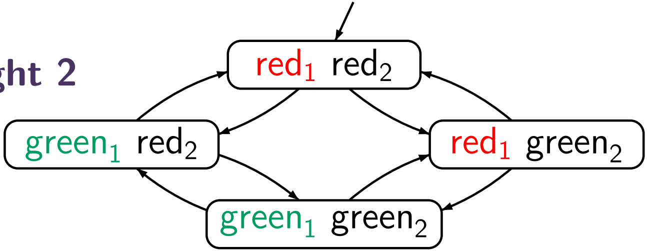

light 1 ||| light 2 $\not\models$ "infinitely often ***green***$_1$"

Two independent traffic lights

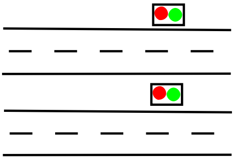

light 1

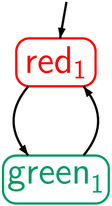

light 2

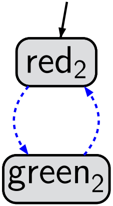

light 1 ||| light 2

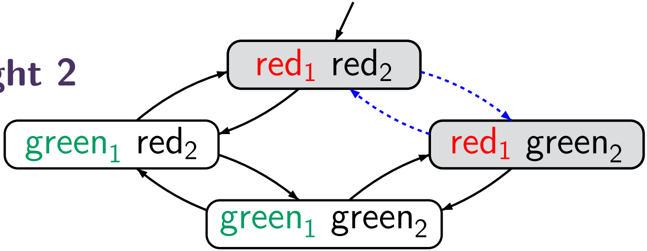

light 1 ||| light 2 $\not\models$ "infinitely often ***$green_1$***"

Two independent traffic lights

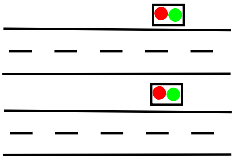

light 1

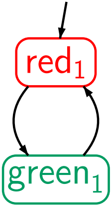

light 2

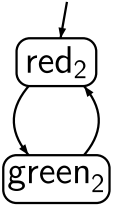

light 1 ||| light 2

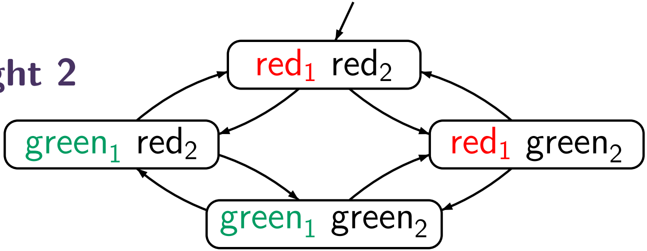

light 1 ||| light 2 $\not\models$ "infinitely often ***green$_1$***"

although **light 1** $\models$ "infinitely often ***green$_1$***"

Two independent traffic lights

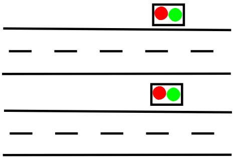

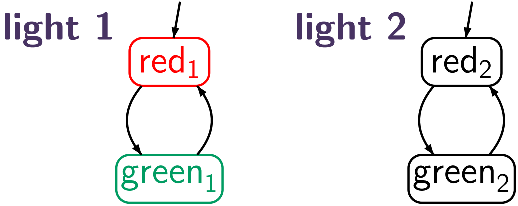

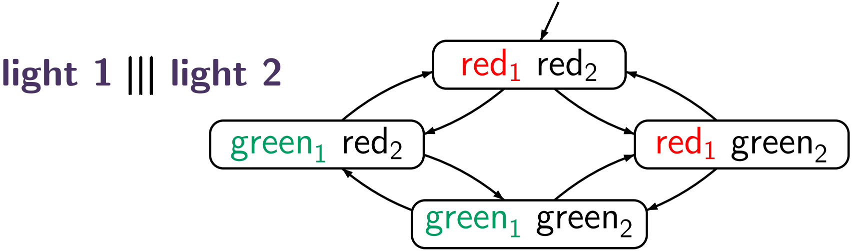

light 1 ||| light 2 $\not\models$ "infinitely often ***green$_1$***"

interleaving is completely time abstract !

Mutual exclusion (semaphore)

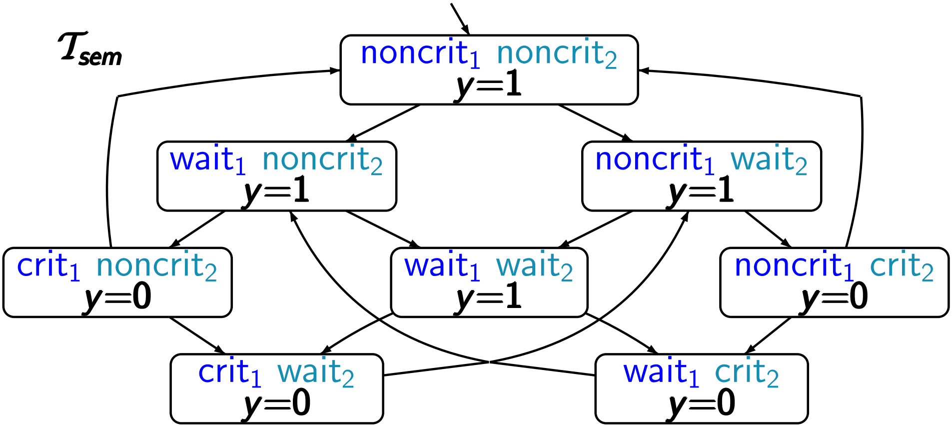

Mutual exclusion (semaphore)

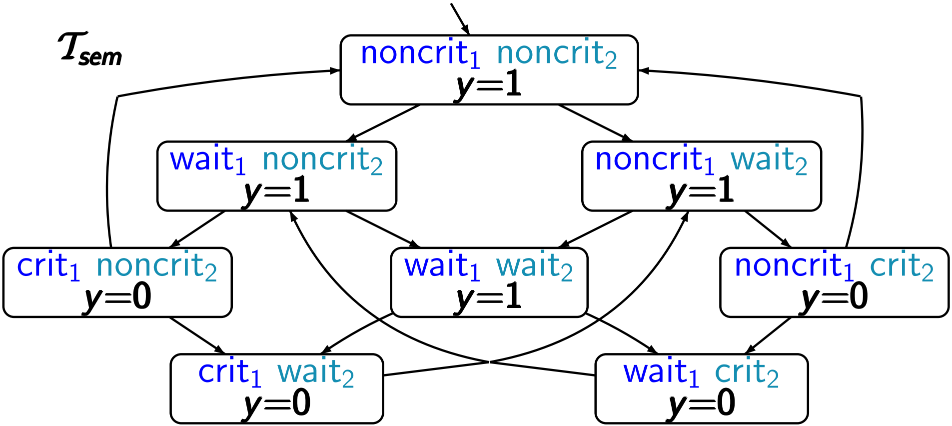

liveness property $\widehat{=}$ "each waiting process will eventually enter its critical section"

Mutual exclusion (semaphore)

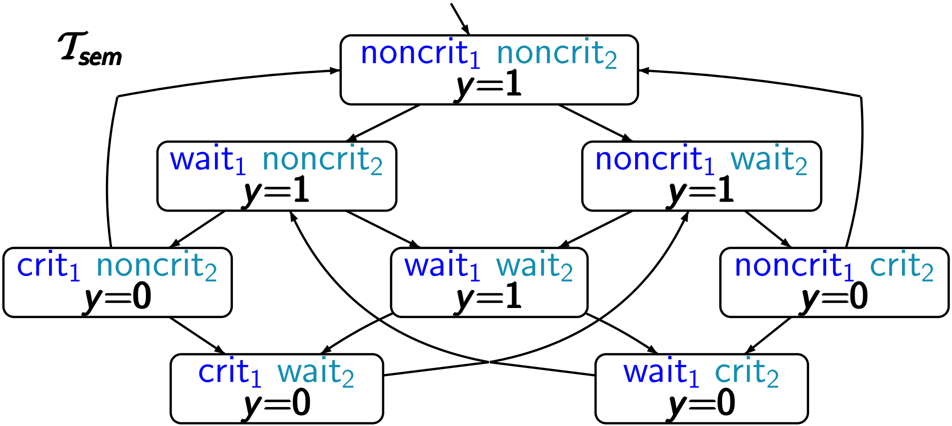

$\mathcal{T}_{sem} \not\models$ "each waiting process will eventually enter its critical section"

Mutual exclusion (semaphore)

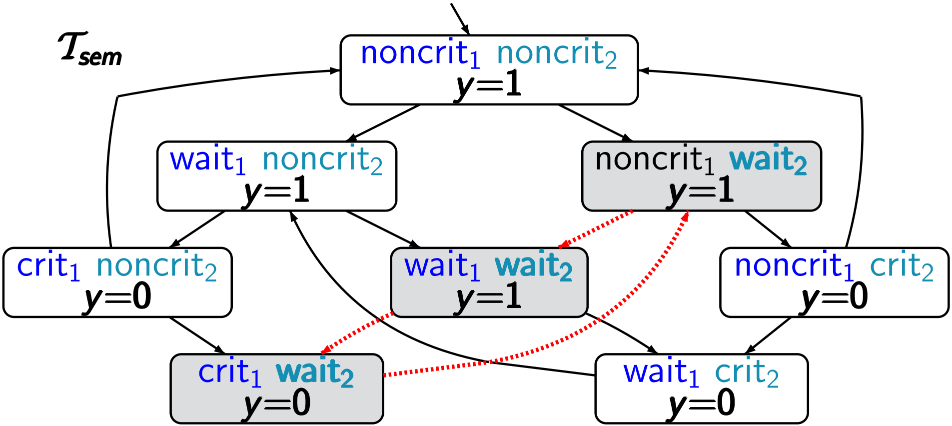

$\mathcal{T}_{sem} \not\models$ "each waiting process will eventually enter its critical section"

Mutual exclusion (semaphore)

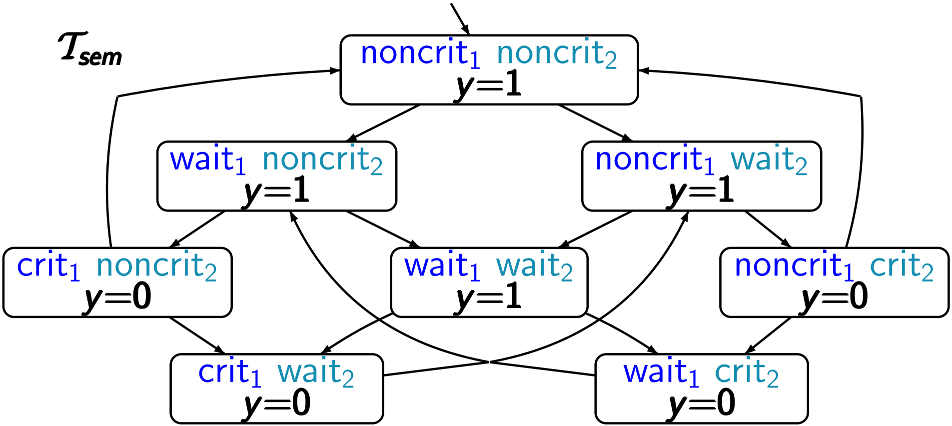

$\mathcal{T}_{sem} \not\models$ "each waiting process will eventually enter its critical section"

level of abstraction is too coarse !

Process fairness

LF2.6-5

Process fairness

two independent non-communicating processes $P_1 \mid\mid\mid P_2$

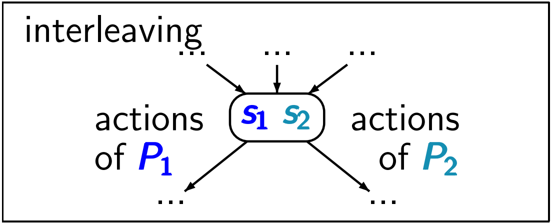

possible interleavings:

P_1 P_2 P_2 P_1 P_1 P_1 P_2 P_1 P_2 P_2 P_2 P_1 P_1 ...
P_1 P_1 P_2 P_1 P_1 P_2 P_1 P_1 P_2 P_1 P_1 P_2 P_1 ...

Process fairness

two independent non-communicating processes $P_1 \mid\mid\mid P_2$

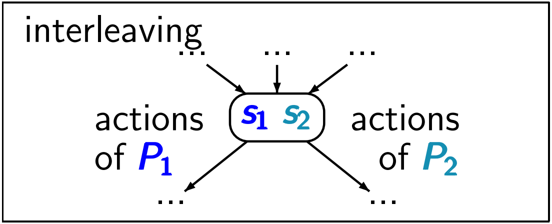

possible interleavings:

$P_1\ P_2\ P_2\ P_1\ P_1\ P_1\ P_2\ P_1\ P_2\ P_2\ P_2\ P_1\ P_1\ \ldots$

$P_1\ P_1\ P_2\ P_1\ P_1\ P_2\ P_1\ P_1\ P_2\ P_1\ P_1\ P_2\ P_1\ \ldots$

$P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ \ldots$

Process fairness

two independent non-communicating processes $P_1 \; ||| \; P_2$

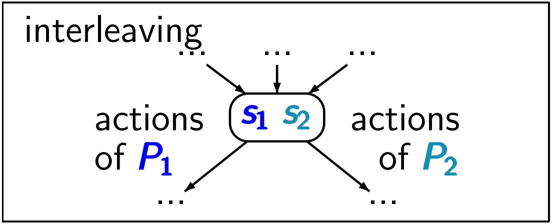

possible interleavings:

$P_1\ P_2\ P_2\ P_1\ P_1\ P_1\ P_2\ P_1\ P_2\ P_2\ P_2\ P_1\ P_1\ \ldots$ fair

$P_1\ P_1\ P_2\ P_1\ P_1\ P_2\ P_1\ P_1\ P_2\ P_1\ P_1\ P_2\ P_1\ \ldots$ fair

$P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ P_1\ \ldots$ unfair

Process fairness

two independent non-communicating processes $P_1 \mid\mid\mid P_2$

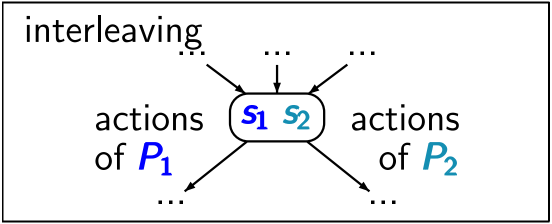

possible interleavings:

P_1 P_2 P_2 P_1 P_1 P_1 P_2 P_1 P_2 P_2 P_2 P_1 P_1 ... fair

P_1 P_1 P_2 P_1 P_1 P_2 P_1 P_1 P_2 P_1 P_1 P_2 P_1 ... fair

P_1 P_1 P_1 P_1 P_1 P_1 P_1 P_1 P_1 P_1 P_1 P_1 P_1 ... unfair

process fairness assumes an appropriate resolution of the nondeterminism resulting from interleaving and competitions

Nuances of fairness

LF2.6-6

- unconditional fairness

- strong fairness

- weak fairness

Nuances of fairness

- unconditional fairness, e.g.,

  every process enters gets its turn infinitely often.

- strong fairness

- weak fairness

Nuances of fairness

- unconditional fairness, e.g.,

  every process enters gets its turn infinitely often.

- strong fairness, e.g.,

  every process that is enabled infinitely often
  gets its turn infinitely often.

- weak fairness

Nuances of fairness

- unconditional fairness, e.g.,

  every process enters gets its turn infinitely often.

- strong fairness, e.g.,

  every process that is enabled infinitely often
  gets its turn infinitely often.

- weak fairness, e.g.,

  every process that is continuously enabled
  from a certain time instance on,
  gets its turn infinitely often.

Fairness for action-set

LF2.6-7

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

we will provide conditions for

- unconditional A-fairness of ρ
- strong A-fairness of ρ
- weak A-fairness of ρ

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

we will provide conditions for

- unconditional A-fairness of ρ
- strong A-fairness of ρ
- weak A-fairness of ρ

using the following notations:

$$Act(s_i) \;=\; \{\beta \in Act : \exists s' \text{ s.t. } s_i \xrightarrow{\beta} s'\}$$

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

we will provide conditions for

- unconditional A-fairness of ρ
- strong A-fairness of ρ
- weak A-fairness of ρ

using the following notations:

$$Act(s_i) = \{\beta \in Act : \exists s' \text{ s.t. } s_i \xrightarrow{\beta} s'\}$$

$$\overset{\infty}{\exists} \ \widehat{=} \ \text{"there exists infinitely many ..."}$$

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

we will provide conditions for

- unconditional A-fairness of ρ
- strong A-fairness of ρ
- weak A-fairness of ρ

using the following notations:

$$Act(s_i) = \{\beta \in Act : \exists s' \text{ s.t. } s_i \xrightarrow{\beta} s'\}$$

$\overset{\infty}{\exists} \;\widehat{=}\;$ "there exists infinitely many ..."

$\overset{\infty}{\forall} \;\widehat{=}\;$ "for all, but finitely many ..."

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and
$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$

$\uparrow$

> "actions in A will be taken infinitely many times"

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and
$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$
- ρ is strongly A-fair, if

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$
- ρ is strongly A-fair, if

$$\overset{\infty}{\exists}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

"If infinitely many times some action in A is enabled, then actions in A will be taken infinitely many times."

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$
- ρ is strongly A-fair, if

$$\overset{\infty}{\exists}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

- ρ is weakly A-fair, if

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$
- ρ is strongly A-fair, if

$$\overset{\infty}{\exists}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

- ρ is weakly A-fair, if

$$\overset{\infty}{\forall}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

"If from some moment, actions in A are enabled, then actions in A will be taken infinitely many times."

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$
- ρ is strongly A-fair, if

$$\overset{\infty}{\exists}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

- ρ is weakly A-fair, if

$$\overset{\infty}{\forall}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

unconditionally A-fair $\Longrightarrow$ strongly A-fair
$\Longrightarrow$ weakly A-fair

Fairness for action-set

Let $\mathcal{T}$ be a TS with action-set Act, $A \subseteq Act$ and

$\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \ldots$ an infinite execution fragment

- ρ is unconditionally A-fair, if $\overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$
- ρ is strongly A-fair, if

$$\overset{\infty}{\exists}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

- ρ is weakly A-fair, if

$$\overset{\infty}{\forall}\, i \geq 0.\, A \cap Act(s_i) \neq \varnothing \quad \Longrightarrow \quad \overset{\infty}{\exists}\, i \geq 0.\, \alpha_i \in A$$

unconditionally A-fair	$\Longrightarrow$	strongly A-fair
	$\Longrightarrow$	weakly A-fair

Strong and weak action fairness

strong ***A***-fairness is *violated* if

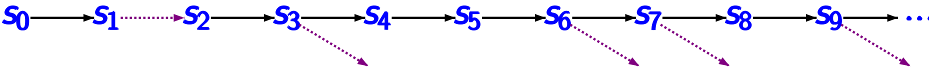

- no ***A***-actions are executed from a certain moment
- ***A***-actions are enabled infinitely many times

Strong and weak action fairness

strong A-fairness is *violated* if

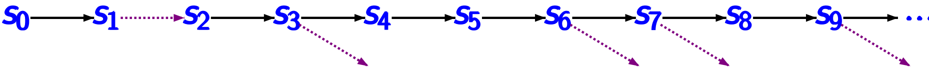

- no A-actions are executed from a certain moment
- A-actions are enabled infinitely many times

weak A-fairness is *violated* if

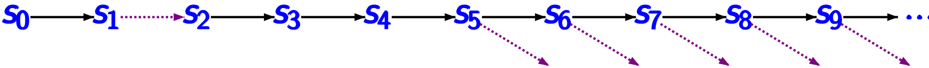

- no A-actions are executed from a certain moment
- A-actions are continuously enabled from some moment on

Mutual exclusion with arbiter

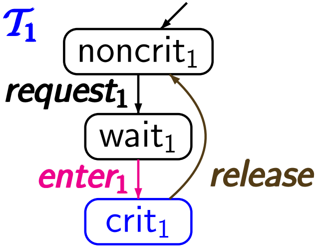

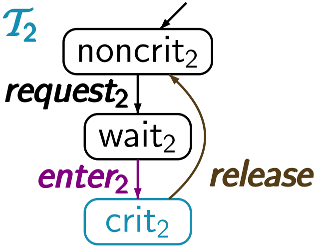

Mutual exclusion with arbiter

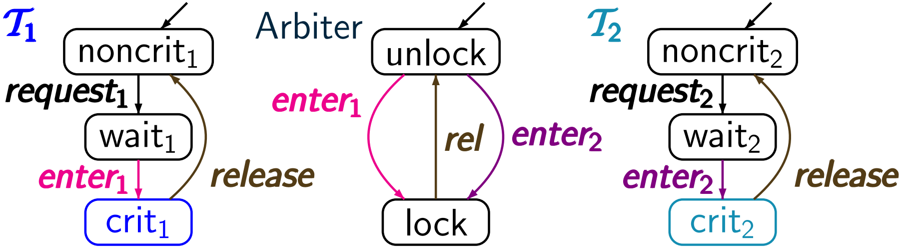

Mutual exclusion with arbiter

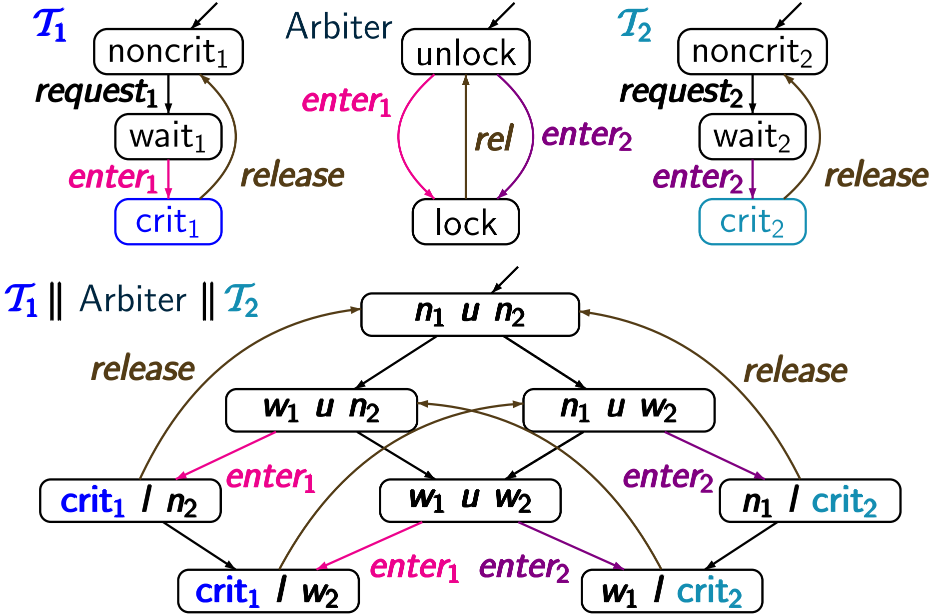

Unconditional, strongly or weakly fair?

$\mathcal{T}_1 \,\|\, \text{Arbiter} \,\|\, \mathcal{T}_2$

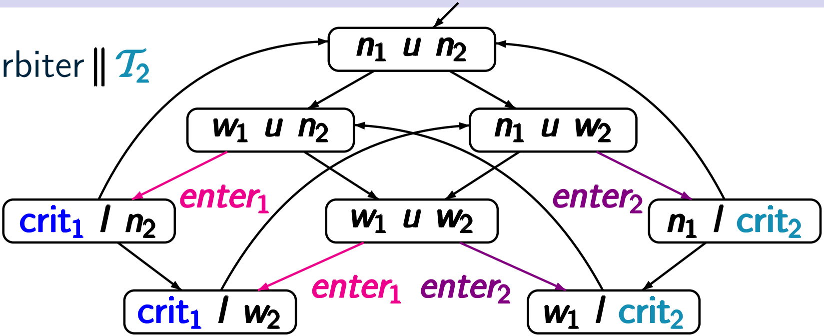

Unconditional, strongly or weakly fair?

$\mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

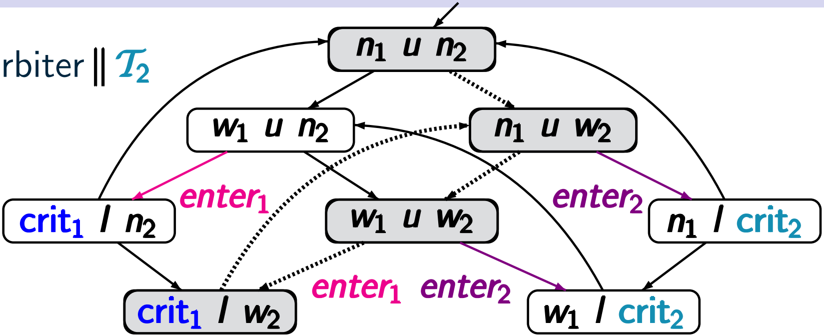

fairness for action set $A = \{enter_1\}$:

$$\langle n_1, u, n_2 \rangle \rightarrow \Big(\langle n_1, u, w_2 \rangle \rightarrow \langle w_1, u, w_2 \rangle \rightarrow \langle crit_1, l, w_2 \rangle \Big)^{\omega}$$

- unconditional A-fairness:
- strong A-fairness:
- weak A-fairness:

Unconditional, strongly or weakly fair?

$\mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

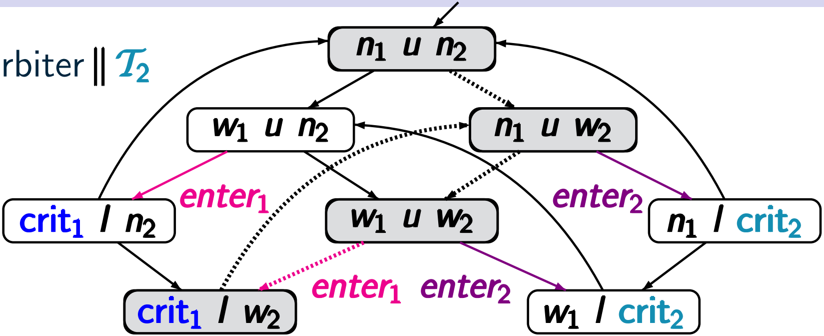

fairness for action set $A = \{enter_1\}$:

$$\langle n_1, u, n_2 \rangle \rightarrow \Big(\langle n_1, u, w_2 \rangle \rightarrow \langle w_1, u, w_2 \rangle \rightarrow \langle crit_1, l, w_2 \rangle \Big)^{\omega}$$

- unconditional A-fairness: **yes**
- strong A-fairness: **yes** ← unconditionally fair
- weak A-fairness: **yes** ← unconditionally fair

Unconditional, strongly or weakly fair?

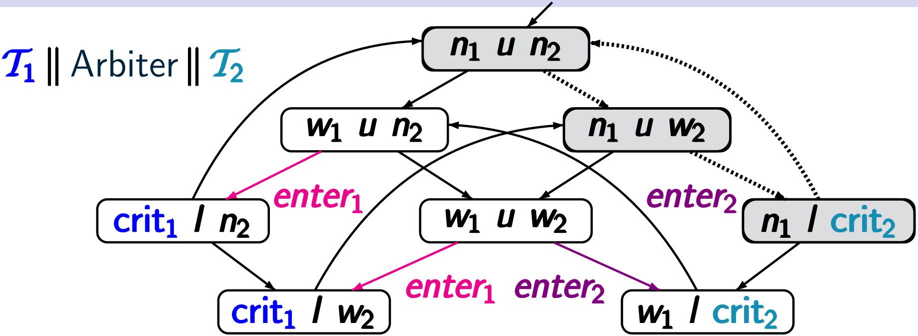

fairness for action-set $A = \{enter_1\}$:

$$\Big(\langle n_1, u, n_2 \rangle \longrightarrow \langle n_1, u, w_2 \rangle \longrightarrow \langle n_1, l, \text{crit}_2 \rangle \Big)^{\omega}$$

- unconditional A-fairness:
- strong A-fairness:
- weak A-fairness:

Unconditional, strongly or weakly fair?

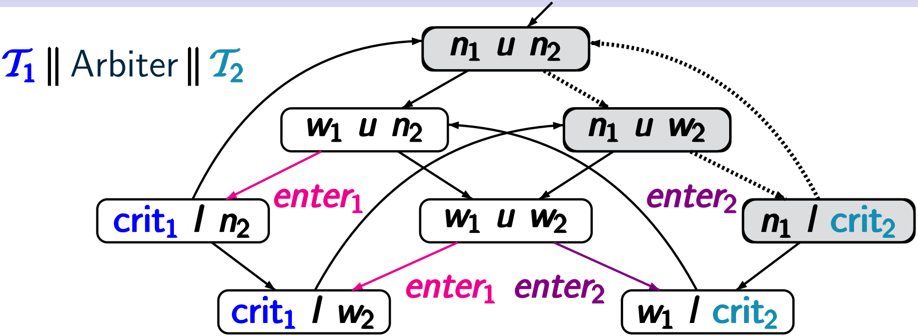

fairness for action-set $A = \{enter_1\}$:

$$\Big(\langle n_1, u, n_2\rangle \longrightarrow \langle n_1, u, w_2\rangle \longrightarrow \langle n_1, l, crit_2\rangle\Big)^\omega$$

- unconditional A-fairness: **no**
- strong A-fairness: **yes** $\leftarrow$ A never enabled
- weak A-fairness: **yes** $\leftarrow$ strongly A-fair

Unconditional, strongly or weakly fair?

$\mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

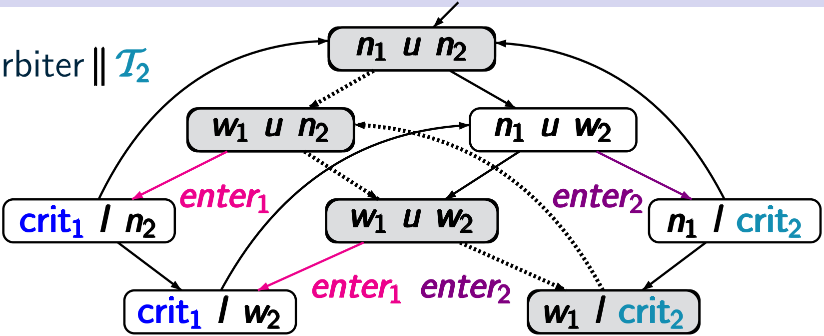

fairness for action-set $A = \{enter_1\}$:

$$\langle n_1, u, n_2\rangle \rightarrow \Big(\langle w_1, u, n_2\rangle \rightarrow \langle w_1, u, w_2\rangle \rightarrow \langle n_1, l, crit_2\rangle\Big)^{\omega}$$

- unconditional A-fairness:
- strong A-fairness:
- weak A-fairness:

Unconditional, strongly or weakly fair?

$\mathcal{T}_1 \,\|\, \text{Arbiter} \,\|\, \mathcal{T}_2$

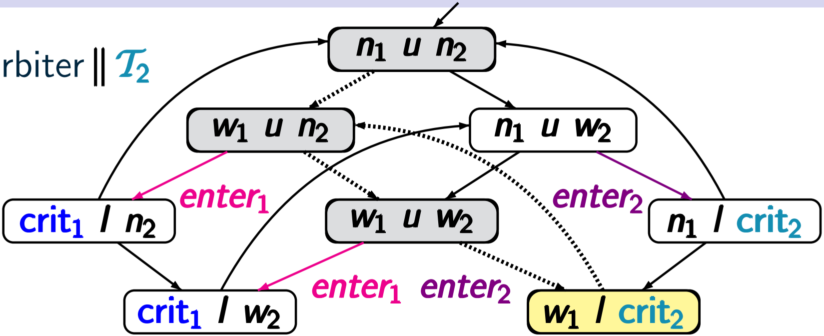

fairness for action-set $A = \{enter_1\}$:

$$\langle n_1, u, n_2\rangle \longrightarrow \Big(\langle w_1, u, n_2\rangle \longrightarrow \langle w_1, u, w_2\rangle \longrightarrow \langle n_1, l, crit_2\rangle \Big)^{\omega}$$

- unconditional A-fairness: **no**
- strong A-fairness: **no**
- weak A-fairness: **yes**

Unconditional, strongly or weakly fair?

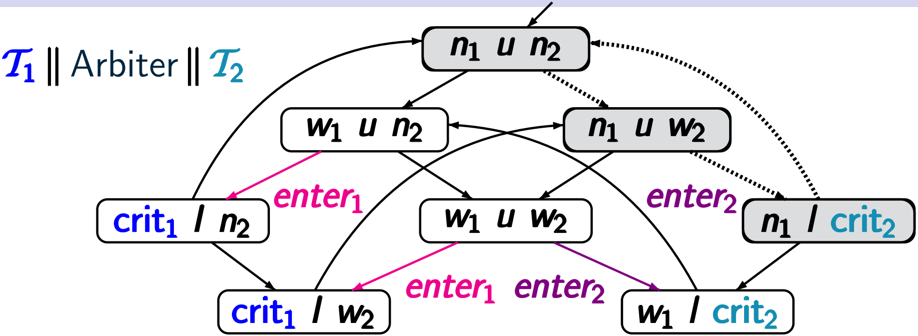

fairness for action set $A = \{enter_1, enter_2\}$:

$$\Big(\langle n_1, u, n_2 \rangle \longrightarrow \langle n_1, u, w_2 \rangle \longrightarrow \langle n_1, u, crit_2 \rangle \Big)^{\omega}$$

- unconditional A-fairness:
- strong A-fairness:
- weak A-fairness:

Unconditional, strongly or weakly fair?

$\mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

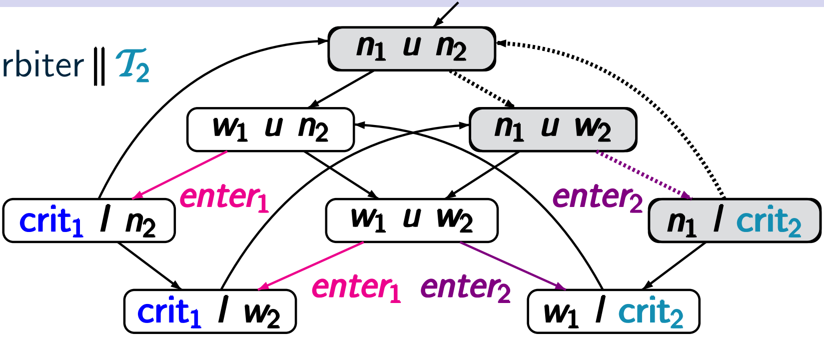

fairness for action set $A = \{enter_1, enter_2\}$:

$$\Big(\langle n_1, u, n_2\rangle \to \langle n_1, u, w_2\rangle \to \langle n_1, u, crit_2\rangle\Big)^{\omega}$$

- unconditional A-fairness: **yes**
- strong A-fairness: **yes**
- weak A-fairness: **yes**

Action-based fairness assumptions

Action-based fairness assumptions

Let $\mathcal{T}$ be a transition system with action-set Act.
A fairness assumption for $\mathcal{T}$ is a triple

$$\mathcal{F} = (\mathcal{F}_{ucond}, \mathcal{F}_{strong}, \mathcal{F}_{weak})$$

where $\mathcal{F}_{ucond}$, $\mathcal{F}_{strong}$, $\mathcal{F}_{weak} \subseteq 2^{Act}$.

Action-based fairness assumptions

Let $\mathcal{T}$ be a transition system with action-set Act.
A fairness assumption for $\mathcal{T}$ is a triple

$$\mathcal{F} = (\mathcal{F}_{ucond}, \mathcal{F}_{strong}, \mathcal{F}_{weak})$$

where $\mathcal{F}_{ucond}$, $\mathcal{F}_{strong}$, $\mathcal{F}_{weak} \subseteq 2^{Act}$.

An execution ρ is called $\mathcal{F}$-fair iff

- ρ is unconditionally A-fair for all $A \in \mathcal{F}_{ucond}$
- ρ is strongly A-fair for all $A \in \mathcal{F}_{strong}$
- ρ is weakly A-fair for all $A \in \mathcal{F}_{weak}$

Action-based fairness assumptions

Let $\mathcal{T}$ be a transition system with action-set Act.
A fairness assumption for $\mathcal{T}$ is a triple

$$\mathcal{F} = (\mathcal{F}_{ucond}, \mathcal{F}_{strong}, \mathcal{F}_{weak})$$

where $\mathcal{F}_{ucond}$, $\mathcal{F}_{strong}$, $\mathcal{F}_{weak} \subseteq 2^{Act}$.

An execution ρ is called $\mathcal{F}$-fair iff

- ρ is unconditionally A-fair for all $A \in \mathcal{F}_{ucond}$
- ρ is strongly A-fair for all $A \in \mathcal{F}_{strong}$
- ρ is weakly A-fair for all $A \in \mathcal{F}_{weak}$

$$FairTraces_{\mathcal{F}}(\mathcal{T}) \stackrel{\text{def}}{=} \{trace(\rho) : \rho \text{ is a } \mathcal{F}\text{-fair execution of } \mathcal{T}\}$$

Fair satisfaction relation

Fair satisfaction relation

A fairness assumption for $\mathcal{T}$ is a triple

$$\mathcal{F} = (\mathcal{F}_{ucond}, \mathcal{F}_{strong}, \mathcal{F}_{weak})$$

where $\mathcal{F}_{ucond}$, $\mathcal{F}_{strong}$, $\mathcal{F}_{weak} \subseteq 2^{Act}$.

An execution ρ is called $\mathcal{F}$-fair iff

- ρ is unconditionally A-fair for all $A \in \mathcal{F}_{ucond}$
- ρ is strongly A-fair for all $A \in \mathcal{F}_{strong}$
- ρ is weakly A-fair for all $A \in \mathcal{F}_{weak}$

If $\mathcal{T}$ is a TS and E a LT property over AP then:

$$\mathcal{T} \models_{\mathcal{F}} E \quad \overset{\text{def}}{\Longleftrightarrow} \quad FairTraces_{\mathcal{F}}(\mathcal{T}) \subseteq E$$

Example: fair satisfaction relation

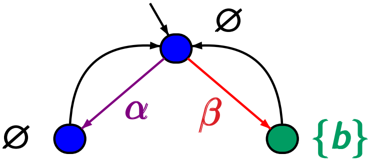

fairness assumption $\mathcal{F}$

- no unconditional fairness condition
- strong fairness for $\{\alpha, \beta\}$
- no weak fairness condition

Example: fair satisfaction relation

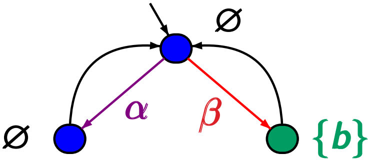

fairness assumption $\mathcal{F}$

- no unconditional fairness condition $\leftarrow \mathcal{F}_{ucond} = \varnothing$
- strong fairness for $\{\alpha, \beta\}$ $\leftarrow \mathcal{F}_{strong} = \{\{\alpha, \beta\}\}$
- no weak fairness condition $\leftarrow \mathcal{F}_{weak} = \varnothing$

Example: fair satisfaction relation

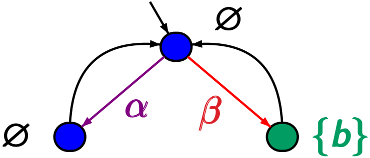

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b" ?

fairness assumption $\mathcal{F}$

- no unconditional fairness condition $\leftarrow \mathcal{F}_{ucond} = \varnothing$
- strong fairness for $\{\alpha, \beta\}$ $\leftarrow \mathcal{F}_{strong} = \{\{\alpha, \beta\}\}$
- no weak fairness condition $\leftarrow \mathcal{F}_{weak} = \varnothing$

Example: fair satisfaction relation

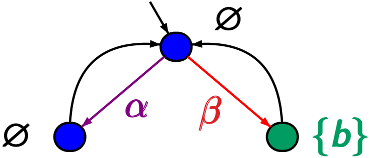

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b" ?

answer: **no**

fairness assumption $\mathcal{F}$

- no unconditional fairness condition $\leftarrow \mathcal{F}_{ucond} = \varnothing$
- strong fairness for $\{\alpha, \beta\}$ $\leftarrow \mathcal{F}_{strong} = \{\{\alpha, \beta\}\}$
- no weak fairness condition $\leftarrow \mathcal{F}_{weak} = \varnothing$

Example: fair satisfaction relation

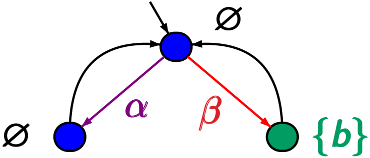

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b" ?

answer: **no**

fairness assumption $\mathcal{F}$

- no unconditional fairness condition $\leftarrow \mathcal{F}_{ucond} = \varnothing$
- strong fairness for $\{\alpha, \beta\}$ $\leftarrow \mathcal{F}_{strong} = \{\{\alpha, \beta\}\}$
- no weak fairness condition $\leftarrow \mathcal{F}_{weak} = \varnothing$

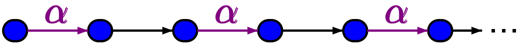

$\mathcal{F}$-fair

actions in $\{\alpha, \beta\}$ are executed infinitely many times

Example: fair satisfaction relation

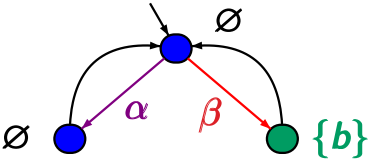

fairness assumption $\mathcal{F}$

- strong fairness for α $\leftarrow \mathcal{F}_{strong} = \{\{\alpha\}\}$
- weak fairness for β $\leftarrow \mathcal{F}_{weak} = \{\{\beta\}\}$
- no unconditional fairness assumption

Example: fair satisfaction relation

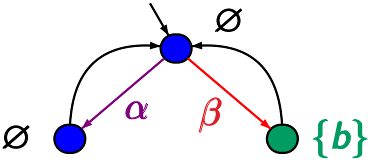

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b" ?

fairness assumption $\mathcal{F}$

- strong fairness for α $\leftarrow \mathcal{F}_{strong} = \{\{\alpha\}\}$
- weak fairness for β $\leftarrow \mathcal{F}_{weak} = \{\{\beta\}\}$
- no unconditional fairness assumption

Example: fair satisfaction relation

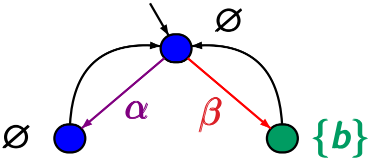

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b" ?

answer: **no**

fairness assumption $\mathcal{F}$

- strong fairness for α $\leftarrow \mathcal{F}_{strong} = \{\{\alpha\}\}$
- weak fairness for β $\leftarrow \mathcal{F}_{weak} = \{\{\beta\}\}$
- no unconditional fairness assumption

Example: fair satisfaction relation

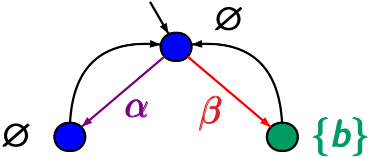

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b" ?

answer: **no**

fairness assumption $\mathcal{F}$

- strong fairness for α $\leftarrow \mathcal{F}_{strong} = \{\{\alpha\}\}$
- weak fairness for β $\leftarrow \mathcal{F}_{weak} = \{\{\beta\}\}$
- no unconditional fairness assumption

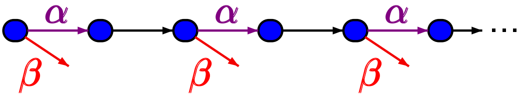

$\mathcal{F}$-fair

Example: fair satisfaction relation

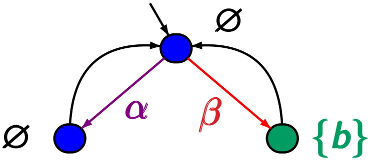

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b"

fairness assumption $\mathcal{F}$

- strong fairness for β $\quad \leftarrow \mathcal{F}_{strong} = \{\{\beta\}\}$
- no weak fairness assumption
- no unconditional fairness assumption

Example: fair satisfaction relation

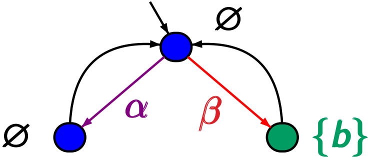

$\mathcal{T} \models_{\mathcal{F}}$ "infinitely often b"

fairness assumption $\mathcal{F}$

- strong fairness for β $\quad\leftarrow \mathcal{F}_{strong} = \{\{\beta\}\}$
- no weak fairness assumption
- no unconditional fairness assumption

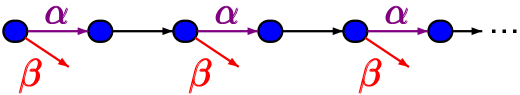

is not $\mathcal{F}$-fair

Which type of fairness?

LF2.6-13A

Which type of fairness?

fairness assumptions should be
as weak as possible

Two independent traffic lights

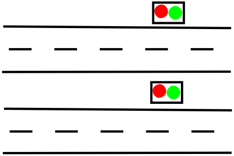

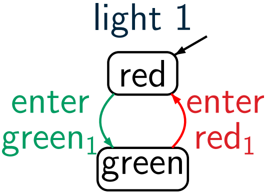

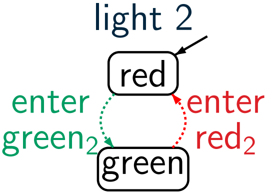

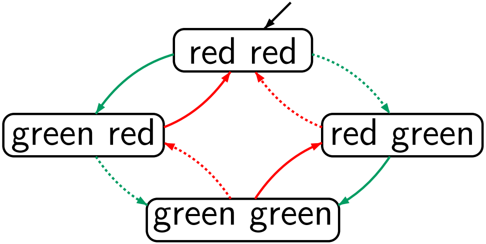

Two independent traffic lights

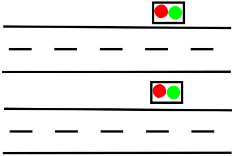

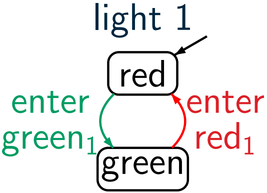

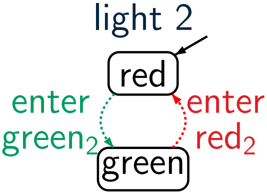

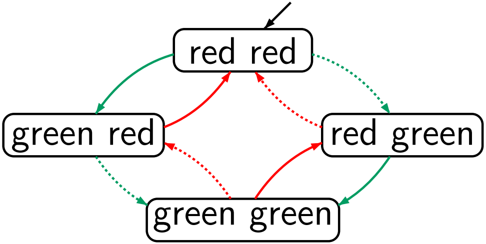

fairness assumption $\mathcal{F}$:

$\mathcal{F}_{ucond} = ?$

$\mathcal{F}_{strong} = ?$

$\mathcal{F}_{weak} = ?$

light 1 ||| light 2 $\models_{\mathcal{F}} E$

$E \mathrel{\widehat{=}}$ "both lights are infinitely often green"

Two independent traffic lights

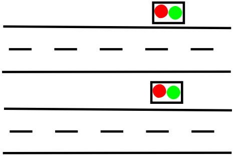

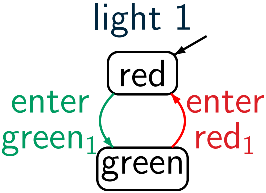

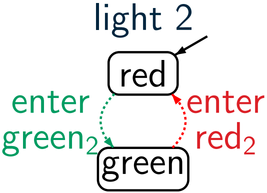

A_1 = actions of light 1
A_2 = actions of light 2

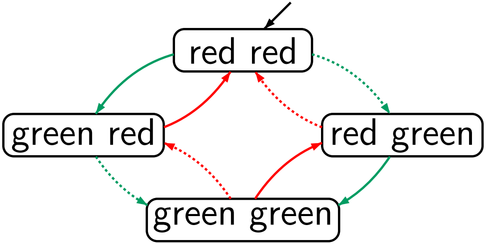

fairness assumption $\mathcal{F}$:

$\mathcal{F}_{ucond} = ?$

$\mathcal{F}_{strong} = ?$

$\mathcal{F}_{weak} = ?$

light 1 $|||$ light 2 $\models_{\mathcal{F}} E$

$E \mathrel{\widehat{=}}$ "both lights are infinitely often green"

Two independent traffic lights

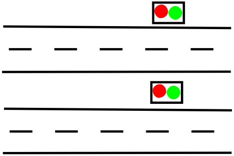

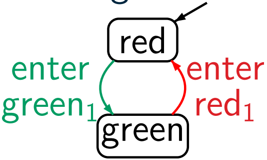

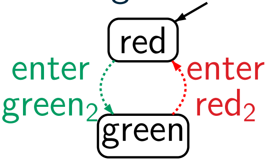

A_1 = actions of light 1

A_2 = actions of light 2

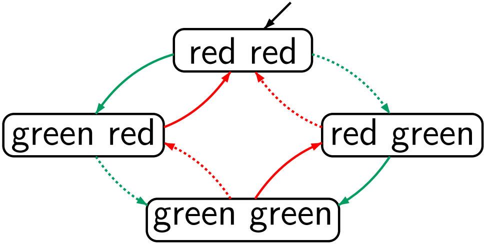

fairness assumption $\mathcal{F}$:

$\mathcal{F}_{ucond} = \varnothing$

$\mathcal{F}_{strong} = \varnothing$

$\mathcal{F}_{weak} = \{A_1, A_2\}$

light 1 $|||$ light 2 $\models_{\mathcal{F}} E$

$E \mathrel{\widehat{=}}$ "both lights are infinitely often green"

Example: MUTEX with fair arbiter

$\mathcal{T} = \mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

Example: MUTEX with fair arbiter

$\mathcal{T} = \mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

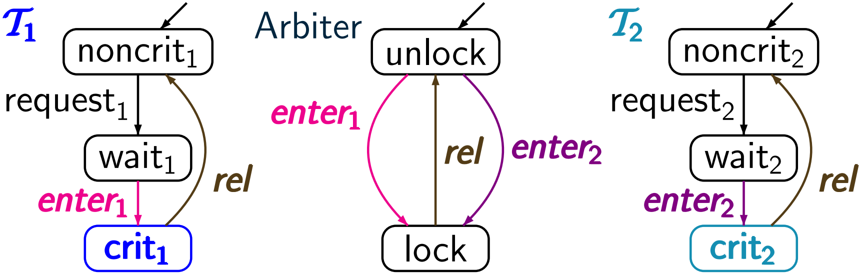

Example: MUTEX with fair arbiter

$\mathcal{T} = \mathcal{T}_1 \parallel \text{Arbiter} \parallel \mathcal{T}_2$

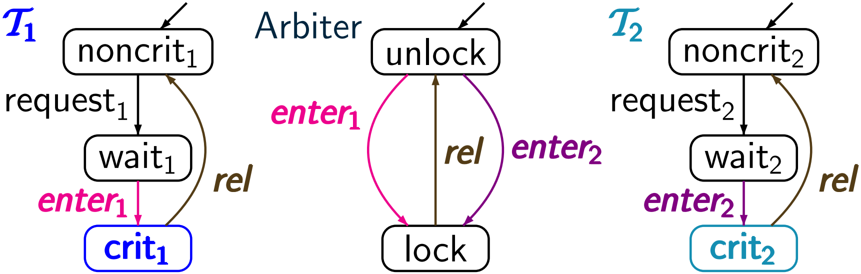

$\mathcal{T}_1$ and $\mathcal{T}_2$ compete to communicate with the arbiter by means of the actions enter_1 and enter_2, respectively

Example: MUTEX with fair arbiter

$\mathcal{T}$

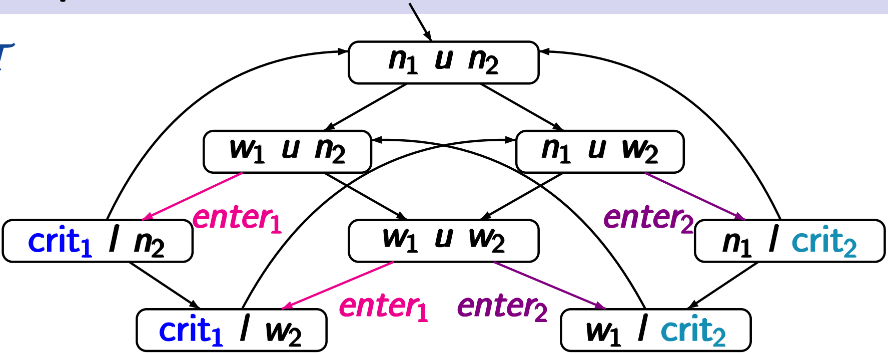

LT property E: each waiting process eventually enters its critical section

$\mathcal{T} \not\models E$

Example: MUTEX with fair arbiter

$\mathcal{T}$

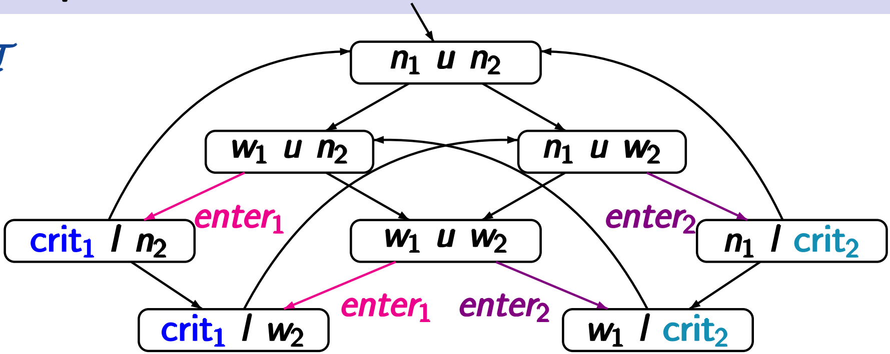

LT property E: each waiting process eventually enters its critical section

fairness assumption $\mathcal{F}$

$\mathcal{F}_{ucond} = \mathcal{F}_{strong} = \varnothing$

$\mathcal{F}_{weak} = \{\{enter_1\}, \{enter_2\}\}$

does $\mathcal{T} \models_{\mathcal{F}} E$ hold ?

Example: MUTEX with fair arbiter

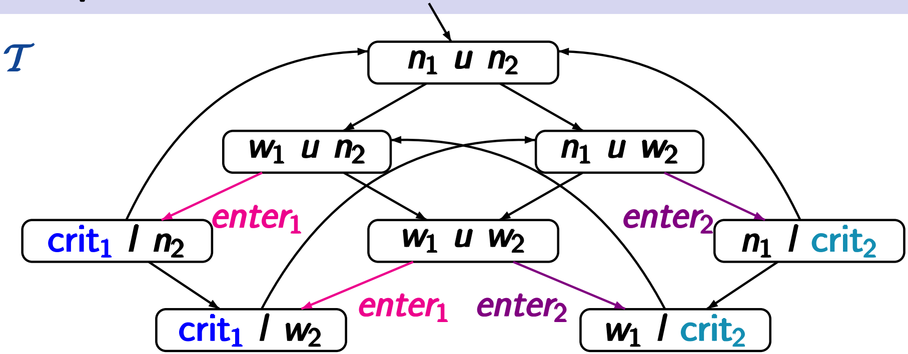

LT property E: each waiting process eventually enters its critical section

fairness assumption $\mathcal{F}$

$\mathcal{F}_{ucond} = \mathcal{F}_{strong} = \varnothing$

$\mathcal{F}_{weak} = \{\{enter_1\}, \{enter_2\}\}$

does $\mathcal{T} \models_{\mathcal{F}} E$ hold ?

answer: **no**

Example: MUTEX with fair arbiter

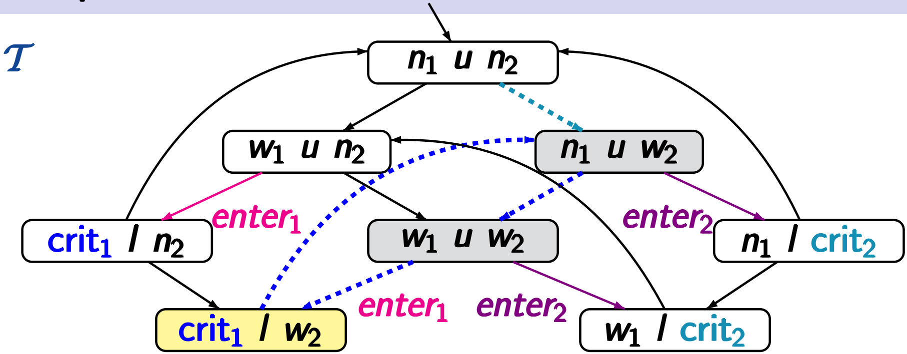

LT property E: each waiting process eventually enters its critical section

fairness assumption $\mathcal{F}$

$\mathcal{F}_{ucond} = \mathcal{F}_{strong} = \varnothing$

$\mathcal{F}_{weak} = \{\{enter_1\}, \{enter_2\}\}$

$\mathcal{T} \not\models_{\mathcal{F}} E$

as $enter_2$ is not enabled in $\langle crit_1, l, w_2 \rangle$

Example: MUTEX with fair arbiter

$\mathcal{T}$

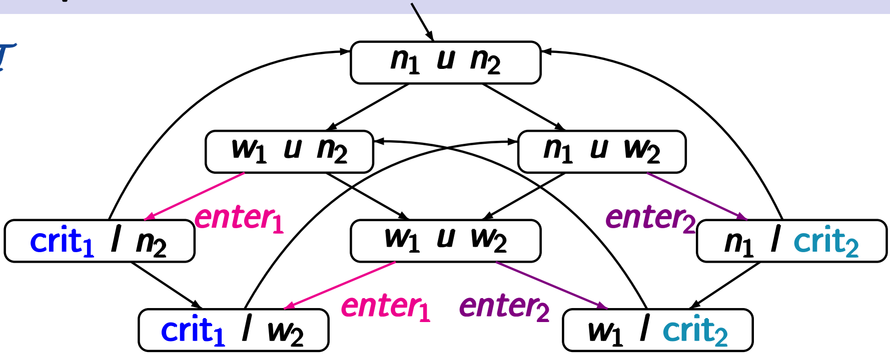

E: each waiting process eventually enters its crit. section

$\mathcal{F}_{ucond}$ = ?
$\mathcal{F}_{strong}$ = ?
$\mathcal{F}_{weak}$ = ?

$\mathcal{T} \not\models E$,

but $\mathcal{T} \models_{\mathcal{F}} E$

Example: MUTEX with fair arbiter

$\mathcal{T}$

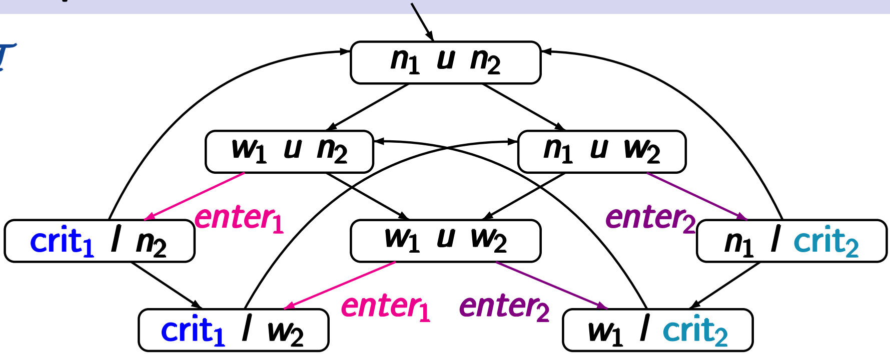

E: each waiting process eventually enters its crit. section

$$\mathcal{F}_{ucond} = \varnothing$$
$$\mathcal{F}_{strong} = \{\{enter_1\}, \{enter_2\}\}$$
$$\mathcal{F}_{weak} = \varnothing$$

$\mathcal{T} \not\models E$,

but $\mathcal{T} \models_{\mathcal{F}} E$

Example: MUTEX with fair arbiter

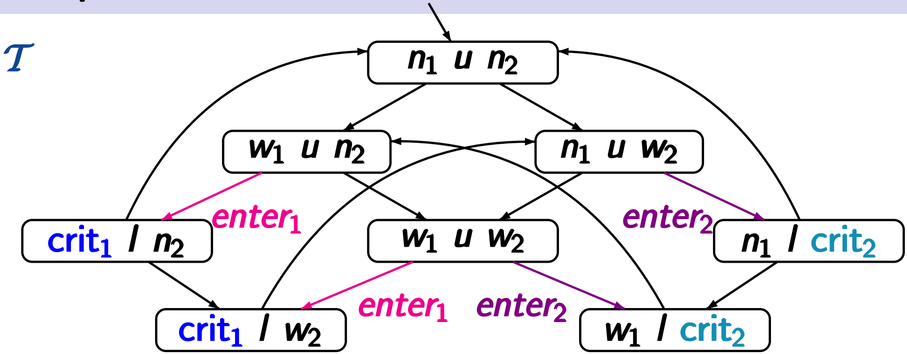

E: each waiting process eventually enters its crit. section

D: each process enters its critical section infinitely often

$$\mathcal{F}_{ucond} = \varnothing$$
$$\mathcal{F}_{strong} = \big\{\{enter_1\}, \{enter_2\}\big\}$$
$$\mathcal{F}_{weak} = \varnothing$$

$$\mathcal{T} \models_{\mathcal{F}} E,$$
$$\mathcal{T} \not\models_{\mathcal{F}} D$$

Example: MUTEX with fair arbiter

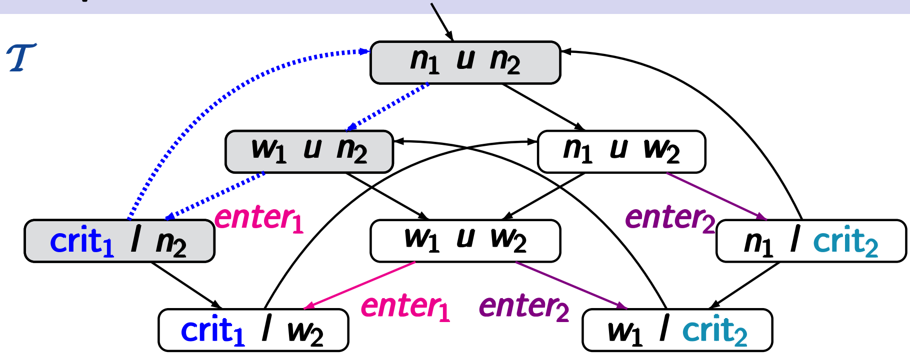

E: each waiting process eventually enters its crit. section

D: each process enters its critical section infinitely often

$$\mathcal{F}_{ucond} = \varnothing$$
$$\mathcal{F}_{strong} = \{\{enter_1\}, \{enter_2\}\}$$
$$\mathcal{F}_{weak} = \varnothing$$

$\mathcal{T} \models_{\mathcal{F}} E$,

$\mathcal{T} \not\models_{\mathcal{F}} D$

Example: MUTEX with fair arbiter

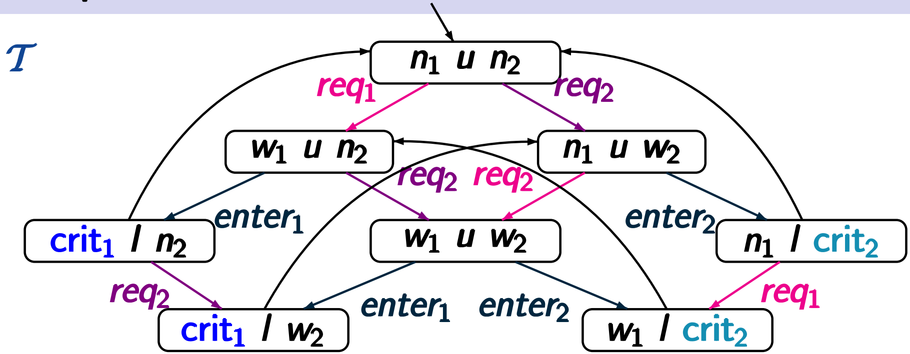

E: each waiting process eventually enters its crit. section

D: each process enters its critical section infinitely often

$$\mathcal{F}_{ucond} = \varnothing$$
$$\mathcal{F}_{strong} = \{\{enter_1\}, \{enter_2\}\}$$
$$\mathcal{F}_{weak} = \{\{req_1\}, \{req_2\}\}$$

$\mathcal{T} \models_{\mathcal{F}} E$,

$\mathcal{T} \models_{\mathcal{F}} D$

Process fairness

LF2.6-19

Process fairness

For asynchronous systems:

parallelism = interleaving + fairness

Process fairness

For asynchronous systems:

$$\text{parallelism} \;=\; \text{interleaving} + \text{fairness}$$

↑ should be as weak as possible

For asynchronous systems:

parallelism = interleaving + fairness

↑ should be as weak as possible

rule of thumb:

- strong fairness for the
  - ∗ choice between dependent actions
  - ∗ resolution of competitions

Process fairness

For asynchronous systems:

parallelism = interleaving + fairness

↑ should be as weak as possible

rule of thumb:

- strong fairness for the
  - ∗ choice between dependent actions
  - ∗ resolution of competitions
- weak fairness for the nondetermism obtained from the interleaving of independent actions

Process fairness

For asynchronous systems:

parallelism = interleaving + fairness

↑ should be as weak as possible

rule of thumb:

- strong fairness for the
  - choice between dependent actions
  - resolution of competitions
- weak fairness for the nondetermism obtained from the interleaving of independent actions
- unconditional fairness: only of theoretical interest

Purpose of fairness conditions

parallelism = interleaving + fairness

Process fairness and other fairness conditions

- can compensate information loss due to interleaving
  or rule out other unrealistic pathological cases
- can be requirements for a scheduler
  or requirements for environment
- can be verifiable system properties

Purpose of fairness conditions

parallelism = interleaving + fairness

Process fairness and other fairness conditions

- can compensate information loss due to interleaving
  or rule out other unrealistic pathological cases
- can be requirements for a scheduler
  or requirements for environment
- can be verifiable system properties

liveness properties:	fairness can be essential
safety properties:	fairness is irrelevant

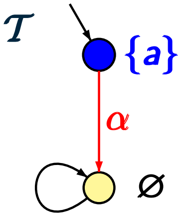

fairness assumption $\mathcal{F}$:
unconditional fairness
for action set $\{\alpha\}$

does $\mathcal{T} \models_{\mathcal{F}}$ "infinitely often a" hold ?

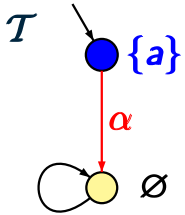

fairness assumption $\mathcal{F}$:
unconditional fairness
for action set $\{\alpha\}$

does $\mathcal{T} \models_{\mathcal{F}}$ "infinitely often *a*" hold ?

answer: **yes** as there is no fair path

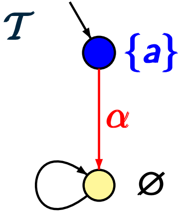

fairness assumption $\mathcal{F}$:
unconditional fairness for action set $\{\alpha\}$

↑ *not* realizable

does $\mathcal{T} \models_{\mathcal{F}}$ "infinitely often a" hold ?

answer: **yes** as there is no fair path

Realizability of fairness assumptions

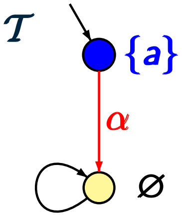

fairness assumption $\mathcal{F}$:
unconditional fairness for action set $\{\alpha\}$

↑

not realizable

does $\mathcal{T} \models_{\mathcal{F}}$ "infinitely often a" hold ?

answer: **yes** as there is no fair path

Realizability requires that each initial finite path fragment can be extended to a $\mathcal{F}$-fair path

Realizability of fairness assumptions

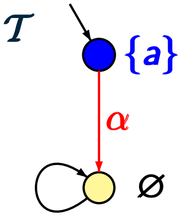

fairness assumption $\mathcal{F}$:
unconditional fairness
for action set $\{\alpha\}$

↑
not realizable

does $\mathcal{T} \models_{\mathcal{F}}$ "infinitely often a" hold ?

answer: **yes** as there is no fair path

Fairness assumption $\mathcal{F}$ is said to be realizable for a transition system $\mathcal{T}$ if for each reachable state s in $\mathcal{T}$ there exists a $\mathcal{F}$-fair path starting in s

Safety and realizable fairness

Safety and realizable fairness

Realizable fairness assumptions are irrelevant for safety properties

Safety and realizable fairness

Realizable fairness assumptions are irrelevant
for safety properties

If $\mathcal{F}$ is a realizable fairness assumption for TS $\mathcal{T}$ and E a safety property then:

$$\mathcal{T} \models E \quad \text{iff} \quad \mathcal{T} \models_{\mathcal{F}} E$$

Safety and realizable fairness

Realizable fairness assumptions are irrelevant for safety properties

If $\mathcal{F}$ is a realizable fairness assumption for TS $\mathcal{T}$ and E a safety property then:

$$\mathcal{T} \models E \quad \text{iff} \quad \mathcal{T} \models_{\mathcal{F}} E$$

... wrong for non-realizable fairness assumptions

Realizable fairness assumptions are irrelevant for safety properties

If $\mathcal{F}$ is a realizable fairness assumption for TS $\mathcal{T}$ and E a safety property then:

$$\mathcal{T} \models E \quad \text{iff} \quad \mathcal{T} \models_{\mathcal{F}} E$$

... wrong for non-realizable fairness assumptions

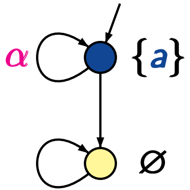

$\mathcal{F}$: unconditional fairness for $\{\alpha\}$

Safety and realizable fairness

Realizable fairness assumptions are irrelevant
for safety properties

> If $\mathcal{F}$ is a realizable fairness assumption for TS $\mathcal{T}$ and E a safety property then:
>
> $$\mathcal{T} \models E \quad \text{iff} \quad \mathcal{T} \models_{\mathcal{F}} E$$

... wrong for non-realizable fairness assumptions

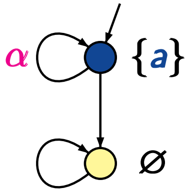

$\mathcal{F}$: unconditional fairness for $\{\alpha\}$

E = invariant "always a"

$\mathcal{T} \not\models E$, but $\mathcal{T} \models_{\mathcal{F}} E$