

Transition Systems and Linear-Time Properties

Lecture #1 of Principles of Model Checking

Joost-Pieter Katoen

Software Modeling and Verification Group

affiliated to University of Twente, Formal Methods and Tools

University of Twente, September 4, 2012

Model checking

- Automated model-based verification and debugging technique
 - model of system = Kripke structure \approx labeled transition system
 - properties expressed in temporal logic like LTL or CTL
 - provides counterexamples in case of property refutation
- Various striking examples
 - Needham-Schroeder security protocol, storm surge barrier, C code
- 2008: Pioneers awarded prestigious ACM Turing Award




Course topics

- Transition systems and linear-time properties
 - traces, safety, liveness, fairness
- Verifying regular linear-time properties
 - omega-regular languages, Büchi automata, nested depth-first search
- LTL model checking
 - syntax, semantics, automata, model-checking algorithm
- CTL model checking
 - syntax, semantics, CTL versus LTL, model-checking algorithm

Course topics

- **Abstraction**
 - stutter (bi)simulation, LTL/CTL equivalence, minimisation algorithms
- **Partial-order reduction**
 - independence, ample sets, dynamic POR
- **Probabilistic model checking**
 - Markov chains, reachability probabilities

Principles of Model Checking



CHRISTEL BAIER

TU Dresden, Germany

JOOST-PIETER KATOEN

RWTH Aachen University, Germany,
and

University of Twente, the Netherlands

"This book offers one of the most comprehensive introductions to logic model checking techniques available today. The authors have found a way to explain both basic concepts and foundational theory thoroughly and in crystal clear prose. Highly recommended for anyone who wants to learn about this important new field, or brush up on their knowledge of the current state of the art."

(Gerard J. Holzmann, NASA JPL, Pasadena)

Content of this lecture

- **Introduction**
 - why model checking?, how to model check?
- **Transition systems**
 - paths, traces, program graphs
- **Linear time properties**
 - safety, liveness, decomposition
- **Fairness**
 - unconditional, strong and weak fairness

Content of this lecture

⇒ Introduction

- why model checking?, how to model check?

● Transition systems

- paths, traces, program graphs


● Linear time properties

- safety, liveness, decomposition

● Fairness

- unconditional, strong and weak fairness

Catching software bugs: the sooner, the better




Milestones in formal verification

- Mathematical approach towards program correctness (Turing, 1949)
- Syntax-based technique for sequential programs (Hoare, 1969)
 - for a given input, does a computer program generate the correct output?
 - based on compositional proof rules expressed in predicate logic
- Syntax-based technique for concurrent programs (Pnueli, 1977)
 - can handle properties referring to situations during the computation
 - based on proof rules expressed in temporal logic
- Automated verification of concurrent programs (Emerson & Clarke, 1981)
 - model-based instead of proof-rule based approach
 - does the concurrent program satisfy a given (logical) property?

these formal techniques are not biased towards the most probable scenarios

Model checking overview



The model checking process

- **Modeling phase**
 - model the system under consideration
 - as a first sanity check, perform some simulations
 - formalise the property to be checked
- **Running phase**
 - run the model checker to check the validity of the property in the model
- **Analysis phase**
 - property satisfied? → check next property (if any)
 - property violated?
 - 1. analyse generated counterexample by simulation
 - 2. refine the model, design, or property . . . and repeat the entire procedure
 - out of memory? → try to reduce the model and try again

Content of this lecture

- **Introduction**
 - why model checking?, how to model check?
- ⇒ **Transition systems**
 - paths, traces, program graphs
- **Linear time properties**
 - safety, liveness, decomposition
- **Fairness**
 - unconditional, strong and weak fairness

Transition systems

- Model to describe the behaviour of systems
- Digraphs where nodes represent *states*, and edges model *transitions*
- **State**:
 - the current colour of a traffic light
 - the current values of all program variables + the program counter
 - the current value of the registers together with the values of the input bits
- **Transition**: (“state change”)
 - a switch from one colour to another
 - the execution of a program statement
 - the change of the registers and output bits for a new input

Formal definition

A *transition system* TS is a tuple $(S, Act, \rightarrow, I, AP, L)$ where

- S is a set of **states**
- Act is a set of **actions**
- $\rightarrow \subseteq S \times Act \times S$ is a **transition relation**
- $I \subseteq S$ is a set of **initial states**
- AP is a set of **atomic propositions**
- $L : S \rightarrow 2^{AP}$ is a **labeling function**

S and Act are either finite or countably infinite

Notation: $s \xrightarrow{\alpha} s'$ instead of $(s, \alpha, s') \in \rightarrow$


Paths

- An *infinite path fragment* π is an infinite state sequence:

$s_0 s_1 s_2 \dots$ such that $s_i \in Post(s_{i-1})$ for all $i > 0$

- Notations for path fragment $\pi = s_0 s_1 s_2 \dots$:
 - $first(\pi) = s_0 = \pi[0]$; let $\pi[j] = s_j$ denote the j -th state of π
 - j -th prefix $\pi[..j] = s_0 s_1 \dots s_j$ and j -th suffix $\pi[j..] = s_j s_{j+1} \dots$
- A *path* of TS is an initial, maximal path fragment
 - a *maximal* path fragment cannot be prolonged
 - a path fragment is *initial* if $s_0 \in I$
- $Paths(s)$ is the set of maximal path fragments π with $first(\pi) = s$

A mutual exclusion algorithm



Traces

- Actions are mainly used to model the (possibility of) interaction
 - synchronous or asynchronous communication
- Here, focus on the states that are visited during executions
 - the states themselves are not “observable”, but just their atomic propositions
- Traces are sequences of the form $L(s_0) L(s_1) L(s_2) \dots$
 - just register the (set of) atomic propositions that are valid along the execution
- For transition systems without terminal states:
 - traces are infinite words over the alphabet 2^{AP} , i.e., they are in $(2^{AP})^\omega$
 - we will (mostly) assume that there are no terminal states

Traces

- Let transition system $TS = (S, Act, \rightarrow, I, AP, L)$ **without terminal states**
- The *trace* of path fragment $\pi = s_0 s_1 \dots$ is $trace(\pi) = L(s_0) L(s_1) \dots$
 - the trace of $\widehat{\pi} = s_0 s_1 \dots s_n$ is $trace(\widehat{\pi}) = L(s_0) L(s_1) \dots L(s_n)$.
- The set of traces of a set Π of paths: $trace(\Pi) = \{ trace(\pi) \mid \pi \in \Pi \}$
- $Traces(s) = trace(Paths(s))$ $Traces(TS) = \bigcup_{s \in I} Traces(s)$

Example traces

Let $AP = \{ crit_1, crit_2 \}$

Example path:

$$\begin{aligned} \pi = & \langle n_1, n_2, y = 1 \rangle \rightarrow \langle w_1, n_2, y = 1 \rangle \rightarrow \langle c_1, n_2, y = 0 \rangle \rightarrow \\ & \langle n_1, n_2, y = 1 \rangle \rightarrow \langle n_1, w_2, y = 1 \rangle \rightarrow \langle n_1, c_2, y = 0 \rangle \rightarrow \dots \end{aligned}$$

The trace of this path is the infinite word:

$$trace(\pi) = \emptyset \emptyset \{ crit_1 \} \emptyset \emptyset \{ crit_2 \} \emptyset \emptyset \{ crit_1 \} \emptyset \emptyset \{ crit_2 \} \dots$$

The trace of the finite path fragment:

$$\begin{aligned} \widehat{\pi} = & \langle n_1, n_2, y = 1 \rangle \rightarrow \langle w_1, n_2, y = 1 \rangle \rightarrow \langle w_1, w_2, y = 1 \rangle \rightarrow \\ & \langle w_1, c_2, y = 0 \rangle \rightarrow \langle w_1, n_2, y = 1 \rangle \rightarrow \langle c_1, n_2, y = 0 \rangle \end{aligned}$$

is:

$$trace(\widehat{\pi}) = \emptyset \emptyset \emptyset \{ crit_2 \} \emptyset \{ crit_1 \}$$

Program graphs: A beverage vending machine

“Abstract” transitions:

$$\begin{array}{lll}
 \text{start} \xrightarrow{\text{true:coin}} \text{select} & \text{and} & \text{start} \xrightarrow{\text{true:refill}} \text{start} \\
 \text{select} \xrightarrow{\text{nsprite}>0:sget} \text{start} & \text{and} & \text{select} \xrightarrow{\text{nbeer}>0:bget} \text{start} \\
 \text{select} \xrightarrow{\text{nsprite}=0 \wedge \text{nbeer}=0:ret_coin} \text{start}
 \end{array}$$

Action	Effect on variables
<i>coin</i>	
<i>ret_coin</i>	
<i>sget</i>	$nsprite := nsprite - 1$
<i>bget</i>	$nbeer := nbeer - 1$
<i>refill</i>	$nsprite := max; nbeer := max$

Some preliminaries

- typed variables with a **valuation** that assigns values to variables
 - e.g., $\eta(x) = 17$ and $\eta(y) = -2$
- the set of Boolean **conditions** over Var
 - propositional logic formulas whose propositions are of the form “ $\overline{x} \in \overline{D}$ ”
 - $(-3 < x \leq 5) \wedge (y = \text{green}) \wedge (x \leq 2 \cdot x')$
- **effect** of the actions is formalized by means of a mapping:

$$Effect : Act \times Eval(Var) \rightarrow Eval(Var)$$

- e.g., $\alpha \equiv x := y + 5$ and evaluation $\eta(x) = 17$ and $\eta(y) = -2$
- $Effect(\alpha, \eta)(x) = \eta(y) + 5 = 3$, and $Effect(\alpha, \eta)(y) = \eta(y) = -2$

Program graphs

A *program graph* PG over set Var of typed variables is a tuple

$$(Loc, Act, Effect, \rightarrow, Loc_0, g_0) \quad \text{where}$$

- Loc is a set of *locations* with initial locations $Loc_0 \subseteq Loc$
- Act is a set of actions
- $Effect : Act \times Eval(Var) \rightarrow Eval(Var)$ is the *effect* function
- $\rightarrow \subseteq Loc \times \underbrace{Cond(Var)}_{\text{Boolean conditions over } Var} \times Act \times Loc$, transition relation
- $g_0 \in Cond(Var)$ is the initial *condition*.

Notation: $\ell \xrightarrow{g:\alpha} \ell'$ denotes $(\ell, g, \alpha, \ell') \in \rightarrow$

Beverage vending machine

- $Loc = \{ start, select \}$ with $Loc_0 = \{ start \}$
- $Act = \{ bget, sget, coin, ret_coin, refill \}$
- $Var = \{ nsprite, nbeer \}$ with domain $\{ 0, 1, \dots, max \}$

$$Effect(coin, \eta) = \eta$$

$$Effect(ret_coin, \eta) = \eta$$

- $Effect(sget, \eta) = \eta[nsprite := nsprite - 1]$

$$Effect(bget, \eta) = \eta[nbeer := nbeer - 1]$$

$$Effect(refill, \eta) = [nsprite := max, nbeer := max]$$

- $g_0 = (nsprite = max \wedge nbeer = max)$

From program graphs to transition systems

- Basic strategy: *unfolding*
 - state = location (current control) ℓ + data valuation η
 - initial state = initial location satisfying the initial condition g_0
- Propositions and labeling
 - propositions: “at ℓ ” and “ $x \in D$ ” for $D \subseteq \text{dom}(x)$
 - $\langle \ell, \eta \rangle$ is labeled with “at ℓ ” and all conditions that hold in η
- $\ell \xrightarrow{g:\alpha} \ell'$ and g holds in η then $\langle \ell, \eta \rangle \xrightarrow{\alpha} \langle \ell', \text{Effect}(\alpha, \eta) \rangle$


Transition systems for program graphs

The transition system $TS(PG)$ of program graph

$$PG = (Loc, \textcolor{blue}{Act}, \textit{Effect}, \longrightarrow, Loc_0, g_0)$$

over set Var of variables is the tuple $(S, \textcolor{blue}{Act}, \longrightarrow, I, AP, L)$ where

- $S = Loc \times \textit{Eval}(Var)$
- $\longrightarrow \subseteq S \times \textit{Act} \times S$ is defined by the rule:
$$\frac{\ell \xrightarrow{\textcolor{red}{g}:\alpha} \ell' \quad \eta \models \textcolor{red}{g}}{\langle \ell, \eta \rangle \xrightarrow{\alpha} \langle \ell', \textit{Effect}(\alpha, \eta) \rangle}$$
- $I = \{ \langle \ell, \eta \rangle \mid \ell \in Loc_0, \eta \models g_0 \}$
- $AP = Loc \cup \textit{Cond}(Var)$ and $L(\langle \ell, \eta \rangle) = \{ \ell \} \cup \{ g \in \textit{Cond}(Var) \mid \eta \models g \}$.



Content of this lecture

- **Introduction**
 - why model checking?, how to model check?
- **Transition systems**
 - paths, traces, program graphs
- ⇒ **Linear time properties**
 - safety, liveness, decomposition
- **Fairness**
 - unconditional, strong and weak fairness

Linear-time properties


- Linear-time properties specify the traces that a TS must exhibit
 - LT-property specifies the admissible behaviour of the system
 - later, a logical formalism will be introduced for specifying LT properties
- A *linear-time property* (LT property) over AP is a subset of $(2^{AP})^\omega$
 - finite words are not needed, as it is assumed that there are no terminal states
- TS (over AP) *satisfies* LT-property P (over AP):

$$TS \models P \quad \text{if and only if} \quad \text{Traces}(TS) \subseteq P$$

- TS satisfies the LT property P if all its “observable” behaviors are admissible


Semaphore-based mutual exclusion

PG_1 :



$y=0$ means “lock is currently possessed”; $y=1$ means “lock is free”

Transition system



How to specify mutual exclusion?

“Always at most one process is in its critical section”

- Let $AP = \{ crit_1, crit_2 \}$
 - other atomic propositions are not of any relevance for this property
- Formalization as LT property

P_{mutex} = set of infinite words $A_0 A_1 A_2 \dots$ with $\{ crit_1, crit_2 \} \not\subseteq A_i$ for all $0 \leq i$


- Contained in P_{mutex} are e.g., the infinite words:

$$\{ crit_1 \} \{ crit_2 \} \{ crit_1 \} \{ crit_2 \} \{ crit_1 \} \{ crit_2 \} \dots \quad \text{and}$$

$$\emptyset \emptyset \emptyset \emptyset \emptyset \emptyset \emptyset \dots$$
 - this does not apply to words of the form: $\{ crit_1 \} \emptyset \{ crit_1, crit_2 \} \dots$

Does the semaphore-based algorithm satisfy P_{mutex} ?

Does the semaphore-based algorithm satisfy P_{mutex} ?



Yes as there is no reachable state labeled with $\{ crit_1, crit_2 \}$

How to specify starvation freedom?

“A process that wants to enter the critical section is eventually able to do so”

- Let $AP = \{ \text{wait}_1, \text{crit}_1, \text{wait}_2, \text{crit}_2 \}$
- Formalization as LT-property

$P_{\text{nostarve}} = \text{set of infinite words } A_0 A_1 A_2 \dots \text{ such that:}$

$$\left(\exists^{\infty} j. \text{wait}_i \in A_j \right) \Rightarrow \left(\exists^{\infty} j. \text{crit}_i \in A_j \right) \quad \text{for each } i \in \{1, 2\}$$

\exists^{∞} stands for “there are infinitely many”.

Does the semaphore-based algorithm satisfy P_{nostarve} ?

No. The trace

$$\begin{aligned} & \emptyset \{ \text{wait}_2 \} \{ \text{wait}_1, \text{wait}_2 \} \{ \text{crit}_1, \text{wait}_2 \} \\ & \{ \text{wait}_2 \} \{ \text{wait}_1, \text{wait}_2 \} \{ \text{crit}_1, \text{wait}_2 \} \dots \end{aligned}$$

is a possible trace of the transition system but not in $P_{nostarve}$

Trace equivalence and LT properties

For TS and TS' be transition systems (over AP):

$$Traces(TS) \subseteq Traces(TS')$$

if and only if

for any LT property P : $TS' \models P$ implies $TS \models P$

$$Traces(TS) = Traces(TS')$$

if and only if

TS and TS' satisfy the same LT properties

Invariants

- LT property P_{inv} over AP is an *invariant* if it has the form:

$$P_{inv} = \{ A_0 A_1 A_2 \dots \in (2^{AP})^\omega \mid \forall j \geq 0. A_j \models \Phi \}$$

- where Φ is a propositional logic formula Φ over AP
- Φ is called an *invariant condition* of P_{inv}

- Note that

$$\begin{aligned} TS \models P_{inv} &\quad \text{iff} \quad \text{trace}(\pi) \in P_{inv} \text{ for all paths } \pi \text{ in } TS \\ &\quad \text{iff} \quad L(s) \models \Phi \text{ for all states } s \text{ that belong to a path of } TS \\ &\quad \text{iff} \quad L(s) \models \Phi \text{ for all states } s \in \text{Reach}(TS) \end{aligned}$$

- Φ has to be fulfilled by all initial states and
 - satisfaction of Φ is invariant under all transitions in the reachable fragment of TS

Safety properties

- Safety properties may impose requirements on finite path fragments
 - and **cannot be verified by considering the reachable states only**
- A safety property which is not an invariant:
 - consider a cash dispenser, also known as automated teller machine (ATM)
 - property “money can only be withdrawn once a correct PIN has been provided”
⇒ not an invariant, since it is not a state property
- But a safety property:
 - **any infinite run violating the property has a finite prefix that is “bad”**
 - i.e., in which money is withdrawn without issuing a PIN before

Safety properties

- LT property P_{safe} over AP is a *safety property* if
 - for all $\sigma \in (2^{AP})^\omega \setminus P_{safe}$ there exists a finite prefix $\widehat{\sigma}$ of σ such that:

$$P_{safe} \cap \left\{ \sigma' \in (2^{AP})^\omega \mid \widehat{\sigma} \text{ is a prefix of } \sigma' \right\} = \emptyset$$
- Path fragment $\widehat{\sigma}$ is a *bad prefix* of P_{safe}
 - let $BadPref(P_{safe})$ denote the set of bad prefixes of P_{safe}
- Path fragment $\widehat{\sigma}$ is a *minimal* bad prefix for P_{safe} :
 - if $\widehat{\sigma} \in BadPref(P_{safe})$ and no proper prefix of $\widehat{\sigma}$ is in $BadPref(P_{safe})$

Example safety properties

Safety properties and finite traces

For transition system TS without terminal states

and safety property P_{safe} :

$TS \models P_{safe}$ if and only if $Traces_{fin}(TS) \cap BadPref(P_{safe}) = \emptyset$

Closure

- For trace $\sigma \in (2^{AP})^\omega$, let $\text{pref}(\sigma)$ be the set of *finite prefixes* of σ :

$$\text{pref}(\sigma) = \{ \hat{\sigma} \in (2^{AP})^* \mid \hat{\sigma} \text{ is a finite prefix of } \sigma \}$$

- if $\sigma = A_0 A_1 \dots$ then $\text{pref}(\sigma) = \{ \varepsilon, A_0, A_0 A_1, A_0 A_1 A_2, \dots \}$
- For property P we have: $\text{pref}(P) = \bigcup_{\sigma \in P} \text{pref}(\sigma)$
- The *closure* of LT property P :

$$\text{closure}(P) = \{ \sigma \in (2^{AP})^\omega \mid \text{pref}(\sigma) \subseteq \text{pref}(P) \}$$

- the set of infinite traces whose finite prefixes are also prefixes of P , or
- infinite traces in the closure of P do not have a prefix that is not a prefix of P

Safety properties and closures

For any LT property P over AP :

P is a safety property if and only if $\text{closure}(P) = P$

Why liveness?

- Safety properties specify that “something bad never happens”
- Doing nothing easily fulfills a safety property
 - as this will never lead to a “bad” situation

⇒ Safety properties are complemented by **liveness** properties

- that require some **progress**

- Liveness properties assert that:
 - “**something good**” will happen eventually [Lamport 1977]

The meaning of liveness



[Lamport 2000]

The question of whether a real system satisfies a liveness property is meaningless; it can be answered only by observing the system for an infinite length of time, and real systems don't run forever.

Liveness is always an approximation to the property we really care about. We want a program to terminate within 100 years, but proving that it does would require addition of distracting timing assumptions.

So, we prove the weaker condition that the program eventually terminates. This doesn't prove that the program will terminate within our lifetimes, but it does demonstrate **the absence of infinite loops**.

Liveness properties

LT property P_{live} over AP is a *liveness* property whenever

$$pref(P_{live}) = (2^{AP})^*$$

- A liveness property is an LT property
 - that *does not rule out any prefix*
- Liveness properties are violated in “infinite time”
 - whereas safety properties are violated in finite time
 - finite traces are of no use to decide whether P holds or not
 - any finite prefix can be extended such that the resulting infinite trace satisfies P

Liveness properties for mutual exclusion

- **Eventually:**
 - each process will eventually enter its critical section
- **Repeated eventually:**
 - each process will enter its critical section infinitely often
- **Starvation freedom:**
 - each waiting process will eventually enter its critical section

Safety vs. liveness

- Are safety and liveness properties disjoint? Yes
- Is any linear-time property a safety or liveness property? No
- But:
for any LT property P an equivalent LT property P' exists
which is a conjunction of a safety and a liveness property

A non-safety and non-liveness property

*“the machine provides infinitely often beer
after initially providing sprite three times in a row”*

- This property consists of **two** parts:
 - it requires beer to be provided infinitely often
⇒ as any finite trace fulfills this, it is a **liveness** property
 - the first three drinks it provides should all be sprite
⇒ bad prefix = one of first three drinks is beer; this is a **safety** property
- Property is thus a conjunction of a safety **and** a liveness property

does this apply to all such properties?


Decomposition theorem

For any LT property P over AP there exists
a safety property P_{safe} and a liveness property P_{live}
(both over AP) such that:

$$P = P_{safe} \cap P_{live}$$

Proposal: $P = \underbrace{closure(P)}_{=P_{safe}} \cap \underbrace{\left(P \cup \left(\left(2^{AP} \right)^\omega \setminus closure(P) \right) \right)}_{=P_{live}}$

Classification of LT properties



Content of this lecture

- **Introduction**
 - why model checking?, how to model check?
- **Transition systems**
 - paths, traces, program graphs
- **Linear time properties**
 - safety, liveness, decomposition

⇒ **Fairness**

- unconditional, strong and weak fairness

Does this program always terminate?


Inc ||| Reset

where


```
proc Inc  =  while < x ≥ 0 do x := x + 1 > od
proc Reset = x := -1
```

x is a shared integer variable that initially has value 0

Is it possible to starve?




Process two starves



Is it fair that process two has infinitely many possibilities to enter the critical section, but never enters it?

Process two starves



Is it fair that process two has infinitely many possibilities to enter the critical section, but only enters it finitely often?

Fairness

- Starvation freedom is often considered under **process fairness**
⇒ there is a fair scheduling of the execution of processes
- **Fairness is typically needed to prove liveness**
 - to prove some form of progress, progress needs to be possible
- Fairness is concerned with a **fair resolution of nondeterminism**
 - such that it is not biased to consistently ignore a possible option
- Problem: liveness properties constrain infinite behaviours
 - but some traces—that are unfair—refute the liveness property

Fairness constraints

- What is wrong with our examples? Nothing!
 - interleaving: not realistic as in no processor is infinitely faster than another
 - semaphore-based mutual exclusion: level of abstraction
- Rule out “unrealistic” executions by imposing *fairness constraints*
 - what to rule out? \Rightarrow different kinds of fairness constraints
- “A process gets its turn infinitely often”
 - always *unconditional fairness*
 - if it is enabled infinitely often *strong fairness*
 - if it is continuously enabled from some point on *weak fairness*

Fairness

This program terminates under unconditional (process) fairness:

```
proc Inc   =  while  $\langle x \geq 0 \text{ do } x := x + 1 \rangle \text{ od}$   
proc Reset =   $x := -1$ 
```

x is a shared integer variable that initially has value 0

Fairness constraints

For $TS = (S, Act, \rightarrow, I, AP, L)$ without terminal states, $A \subseteq Act$,
and infinite execution fragment $\rho = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots$ of TS :

1. ρ is *unconditionally A-fair* whenever: true $\implies \underbrace{\forall k \geq 0. \exists j \geq k. \alpha_j \in A}_{\text{infinitely often } A \text{ is taken}}$
2. ρ is *strongly A-fair* whenever:


$$\underbrace{(\forall k \geq 0. \exists j \geq k. Act(s_j) \cap A \neq \emptyset)}_{\text{infinitely often } A \text{ is enabled}} \implies \underbrace{\forall k \geq 0. \exists j \geq k. \alpha_j \in A}_{\text{infinitely often } A \text{ is taken}}$$

3. ρ is *weakly A-fair* whenever:

$$\underbrace{(\exists k \geq 0. \forall j \geq k. Act(s_j) \cap A \neq \emptyset)}_{A \text{ is eventually always enabled}} \implies \underbrace{\forall k \geq 0. \exists j \geq k. \alpha_j \in A}_{\text{infinitely often } A \text{ is taken}}$$

where $Act(s) = \{\alpha \in Act \mid \exists s' \in S. s \xrightarrow{\alpha} s'\}$

Example (un)fair executions



Which fairness notion to use?

- Fairness constraints aim to rule out “unreasonable” runs
- **Too strong?** \Rightarrow relevant computations ruled out
 - verification yields:
 - “**false**”: error found
 - “**true**”: don’t know as some relevant execution may refute it
- **Too weak?** \Rightarrow too many computations considered
 - verification yields:
 - “**true**”: property holds
 - “**false**”: don’t know, as refutation maybe due to some unreasonable run

often a combination of several fairness constraints is used

Fairness assumptions

- A *fairness assumption* for Act is a triple


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

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

- Execution ρ is \mathcal{F} -fair if:
 - it is unconditionally A -fair **for all** $A \in \mathcal{F}_{ucond}$, and
 - it is strongly A -fair **for all** $A \in \mathcal{F}_{strong}$, and
 - it is weakly A -fair **for all** $A \in \mathcal{F}_{weak}$


fairness assumption $(\emptyset, \mathcal{F}', \emptyset)$ denotes strong fairness; $(\emptyset, \emptyset, \mathcal{F}')$ weak, etc.

Fairness for mutual exclusion




$$\mathcal{F} = (\emptyset, \underbrace{\{ \{ enter_1, enter_2 \} \}}_{\mathcal{F}_{strong}}, \emptyset)$$

Fairness for mutual exclusion



$$\mathcal{F}' = (\emptyset, \underbrace{\left\{ \{ \text{enter}_1 \}, \{ \text{enter}_2 \} \right\}}_{\mathcal{F}_{\text{strong}}}, \emptyset)$$

Fairness for mutual exclusion



$$\mathcal{F}'' = \left(\emptyset, \underbrace{\left\{ \{ enter_1 \}, \{ enter_2 \} \right\}}_{\mathcal{F}_{strong}}, \underbrace{\left\{ \{ req_1 \}, \{ req_2 \} \right\}}_{\mathcal{F}_{weak}} \right)$$

in any \mathcal{F}'' -fair execution each process infinitely often requests access

Fair paths and traces

- Path $s_0 \rightarrow s_1 \rightarrow s_2 \dots$ is **\mathcal{F} -fair** if
 - there exists an \mathcal{F} -fair execution $s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} s_2 \dots$
 - $\text{FairPaths}_{\mathcal{F}}(s)$ denotes the set of \mathcal{F} -fair paths that start in s
 - $\text{FairPaths}_{\mathcal{F}}(TS) = \bigcup_{s \in I} \text{FairPaths}_{\mathcal{F}}(s)$
- Trace σ is **\mathcal{F} -fair** if there exists an \mathcal{F} -fair path π with $\text{trace}(\pi) = \sigma$
 - $\text{FairTraces}_{\mathcal{F}}(s) = \text{trace}(\text{FairPaths}_{\mathcal{F}}(s))$
 - $\text{FairTraces}_{\mathcal{F}}(TS) = \text{trace}(\text{FairPaths}_{\mathcal{F}}(TS))$

Fair satisfaction

- TS *satisfies* LT-property P :


$$TS \models P \quad \text{if and only if} \quad \text{Traces}(TS) \subseteq P$$

- TS satisfies the LT property P if *all* its observable behaviors are admissible
- TS *fairly satisfies* LT-property P wrt. fairness assumption \mathcal{F} :

$$TS \models_{\mathcal{F}} P \quad \text{if and only if} \quad \text{FairTraces}_{\mathcal{F}}(TS) \subseteq P$$

- if all paths in TS are \mathcal{F} -fair, then $TS \models_{\mathcal{F}} P$ if and only if $TS \models P$
- if some path in TS is not \mathcal{F} -fair, then possibly $TS \models_{\mathcal{F}} P$ but $TS \not\models P$

Fairness for mutual exclusion



$TS \not\models$ “every process enters its critical section infinitely often”

and $TS \not\models_{\mathcal{F}'} \text{“every ... often”}$

but $TS \models_{\mathcal{F}''} \text{“every ... often”}$

Fairness and safety properties

For TS and safety property P_{safe} (both over AP)
such that for any $s \in \text{Reach}(TS)$: $\text{FairPaths}_{\mathcal{F}}(s) \neq \emptyset$:

$$TS \models P_{safe} \quad \text{if and only if} \quad TS \models_{\mathcal{F}} P_{safe}$$

Safety properties are thus preserved by “realizable” fairness assumptions