Formal Methods in Software Development

O

Ivano Salvo and Igor Melatti

Computer Science Department



SAPIENZA Università di Roma

Lesson **6**, October 29th, 2019

Lesson 6a:

The Fairness problem



- System models often **abstracts from details** such as, for example, **scheduler policies**.
- Interleaving semantics does not rule out unrealistic behaviour, for example, those in which some processes do not make any progress.

Of course, one can embody a fair process scheduling in the model, as in the case of Peterson mutual exclusion algorithm.

Alternatively, one can **assume some fairness properties**, and perform model checking under such assumptions.

Example: (un)fair Schedulers

Mutual Exclusion with shared variables: The starvation freedom property:

"Once access is requested, a process does not have to wait infinitely long before acquiring access to its critical section"

is violated, just because, **abstracting from the scheduling policy**, a legal execution of the system assignes the critical resource always to the same process.

The property:

"Each of the processes is infinitely often in its critical section"

is violated also by the Peterson protocol, as it **does not exclude** that **a process would never** (or finitely often) **request** to enter its critical section.

Example: Interleaving Semantics



Two independent traffic lights (lesson **1**): Interleaving semantics allows infinite executions in which **only the first traffic light commute**: {red₁, red₂}{green1, red₂} {red₁, red₂}{green1, red₂} {red₁, red₂}{green1, red₂}...

Fairness notions

Unconditional Fairness: "every process gets its turn infinitely often" (without conditions, aka **impartiality**)

Strong Fairness: "every process that is **enabled infinitely often** gets its turn infinitely often" (aka **compassion**)

Weak Fairness: every process that is **continously enabled from a certain point on** gets its turn infinitely often" (aka **justice**)

Many other fairness notions have been introduced in literature and there is no clear consensus about which notion should be used in some scenario.

Fairness def. (action based)

Definition 3.43. Unconditional, Strong, and Weak Fairness

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

equivalent formulas

GFφ

1. ρ is unconditionally A-fair whenever $\stackrel{\infty}{\exists} j. \alpha_j \in A$.

$$\begin{pmatrix} \stackrel{\infty}{\exists} j. Act(s_j) \cap A \neq \emptyset \end{pmatrix} \implies \begin{pmatrix} \stackrel{\infty}{\exists} j. \alpha_j \in A \end{pmatrix}.$$

G F $\varphi \rightarrow$ **G F** ψ

3. ρ is weakly A-fair whenever

2. ρ is strongly A-fair whenever

$$\begin{pmatrix} \stackrel{\infty}{\forall} j. Act(s_j) \cap A \neq \emptyset \end{pmatrix} \implies \begin{pmatrix} \stackrel{\infty}{\exists} j. \alpha_j \in A \end{pmatrix}.$$

FG $\varphi \rightarrow$ **GF** ψ

Here, $\stackrel{\infty}{\exists} j$ stands for "there are infinitely many j" and $\stackrel{\infty}{\forall} j$ for "for nearly all j" in the sense of "for all, except for finitely many j". The variable j, of course, ranges over the natural numbers.

Ex: Shared variables program

proc lnc = while $\langle x \ge 0 \text{ do } x := x + 1 \rangle \text{ od}$

```
proc Reset = x := -1
```

This process terminates only if **unconditional fairness** is assumed...If process Inc or process Reset can execute infinitely often, the concurrent program does not terminate.

```
[Brackets (...) means "atomic actions"]
```

Which notion of fairness we should use? No answer!

Keep in mind: if the fairness constraints are **too strong**, **relevant computation can be ruled out**. By contrast, if the fairness constraints are **too weak**, we refute a property because we consider **unrealistic behaviour** of a system.

Uncond. Fairnes $A \Rightarrow$ Strong Fairness $A \Rightarrow$ Weak Fairness A

Mutual Exclusion Reloaded

The dashed execution fragment **is strongly fair** premises are vacously true), but **not unconditionally fair** for {enter₂}. The dotted **is weakly fair**, **but not strongly fair**. Process 1 requests access infinitely often, but not continously.



Example: How to model Fairness

The strong fairness assumption {{**enter**₁, **enter**₂}} ensure only that one of the two process enter its critical section infinitely often. Probably, {{**enter**₁}, {**enter**₂}} is the intended solution!



Fairness: Linear Time Properties

Definition: Let $P \subseteq (2^{AP})^{\omega}$ be an LT property over *AP* and let \mathcal{F} be a fairness assumption over *A*. A transition system \mathcal{M} **fairly satisfies** *P*, notation $\mathcal{M} \models_{\mathcal{F}} P$, if and only if fairTraces(\mathcal{M}) $\subseteq P$.

If all executions of \mathcal{M} satisfies \mathcal{F} , then $\mathcal{M} \vDash_{\mathcal{F}} P$ iff $\mathcal{M} \vDash P$. More in general, we have that $\mathcal{M} \vDash P \Longrightarrow \mathcal{M} \vDash_{\mathcal{F}} P$.

As said before, we also have:

$$\mathcal{M} \vDash_{\operatorname{weak} \mathcal{F}} P \Longrightarrow \mathcal{M} \vDash_{\operatorname{strong} \mathcal{F}} P \Longrightarrow \mathcal{M} \vDash_{\operatorname{uncond} \mathcal{F}} P$$

Example: Independent Traffic Lights. The fair assumption:

{{switchToGreen₁, switchToRed₁},{switchToGreen₂, switchToRed₂}}

rules out unrealistic behaviour, no matter if this is interpreted as strong, weak or unconditional fairness constraint.

 $\operatorname{TrLight}_1 \parallel \operatorname{TrLight}_2 \vDash_{\mathcal{F}} \mathbf{F} \mathbf{G}$ green \equiv "each traffic light is green infin. often"

Example: Mutual Exclusion

Let us consider again the semaphore based mutual exclusion protocol. Let us define the following fairness constraints:

 $\mathcal{F}_{\text{weak}} = \{\{\text{req}_1\}, \{\text{req}_2\}\} \quad \mathcal{F}_{\text{strong}} = \{\{\text{enter}_1\}, \{\text{enter}_2\}\} \quad \mathcal{F}_{\text{uncond}} = \emptyset$

The strong fairness assumption \mathcal{F}_{strong} does not forbid a process to never release its critical section.

The weak fairness assumption $\mathcal{F}_{\text{weak}}$ implies that a process requires to enter critical section infinitely often (and hence it has to leave infinitely often its critical section).

Also Peterson's protocol ensure that process will enter its critical section if it requires it infinitely often. But it does not ensure that processes leave their critical section. To ensure this, we should impose the weak fairness assumption $\mathcal{F}_{weak} = \{\{req_1\}, \{req_2\}\}.$

Weak or Strong Fairness

Rule of Thumb:

Strong fairness is appropriate to obtain an adequate resolution of **contentions between processes**

Weak fairness suffices for sets of actions that represent the concurrent execution of independent actions (interleaving)

Example: Let us assume we have *n* processes represented by transition systems $\mathcal{M}_i = (S_i, A_i, \rightarrow_i, AP_i, L_i)$ and consider the parallel composition $\mathcal{M}_i = \mathcal{M}_1 \parallel \mathcal{M}_2 \parallel \dots \parallel \mathcal{M}_n$. Let us suppose that each pair of processes synchronize on a set of actions $H_{i,i}$.

Example: Fair Synchronization

The **strong fairness assumption** $\{A_1, A_2, ..., A_n\}$ means that each process progress infinitely often (provided that infinitely often a process has an enabled action to execute). This assumption is satisfied, however, only with internal actions and no sync!

{{ α } | $\alpha \in H_{i,j}$ 0 < *i* < *j* ≤ *n* } forces every sync action to be performed infinitely often.

 ${H_{i,j} \mid 0 < i < j \le n}$ forces every pair of processes to synchronize infinitely often, maybe on the same action.

 $\{\bigcup_{0 \le i \le j \le n} H_{i,j}\}$ just requires that there are infinite synchronization actions, regardless of which are processes involved.

For **internal actions**, the **weak fairness assumption**: $\{A_1 \setminus H_1, ..., A_n \setminus H_n\}$, where $H_i = \bigcup_{i \neq j} H_j$, is appropriate, since internal actions are continously ready to be executed.

Lesson 6b:

Fairness in LTL Model Checking

Fairness is expressible in LTL

The three notions of fairness we have considered can be expressed by LTL formulas of the shape:

Unconditional fairness: **G F** φ Strong fairness: **G F** $\varphi \rightarrow$ **G F** ψ Weak fairness: **F G** $\varphi \rightarrow$ **G F** ψ

The only problem is that LTL formulas are built on atomic propositions that label states: therefore φ and ψ depend on the state labeling and they single out set of states of a transition system \mathcal{M} , i.e. { $s \mid \mathcal{M}, s \vDash \varphi$ }.

However, this is not a limitation...

Mutual Exclusion

The strong (action based) fairness assumption \mathcal{F}_{strong} = {{enter₁}, {enter₂}} can be represented by the (state based) LTL formula (observe that enter₁ can be executed only if P₁ is in the state wait₁ and P₂ is not in its critical section):

sfair₁ = **G F** (wait₁ $\land \neg$ crit₂) \rightarrow **G F** crit₁

The assumption sfair₂ can be defined analogously.

 \mathcal{F}_{strong} does not forbid a process to never leave its critical section. The (action based) weak assumption = {{req₁}, {req₂}} can be encoded as the (state based) LTL formula (observe that the action req₁ is executable only if P₁ is in the state noncrit₁)

wfair₁ = **F G** noncrit₁ \rightarrow **G F** wait₁

The assumption wfair₂ can be defined analogously.

Action vs State based Fairness 1

Kripke structures have no action labels. One can always keep **information into states**.

Let \mathcal{M} be a transition system $(S, A, \rightarrow, I, AP, L)$. We can define the system $\mathcal{M}'=(S', A', \rightarrow', I', AP', L')$ where:

- $A' = A \cup \{ \text{ begin } \},$
- $I' = I \times \{ \text{ begin } \},$
- $S' = I' \cup (S \times A),$
- If $s_0 \rightarrow_{\alpha} s \ (s_0 \in I)$, then $(s_0, \text{ begin}) \rightarrow'_{\alpha} (s, \alpha)$. If $s \rightarrow_{\alpha} s'$ then $(s, \beta) \rightarrow'_{\alpha} (s', \alpha)$
- $AP' = AP \cup \{enabled(\alpha), taken(\alpha) \mid \alpha \in A\}$
- $L'(s, \alpha) = L(s) \cup \{ taken(\alpha) \} \cup \{ enabled(\beta) \mid \beta \in Act(s) \} and L'(s_0, begin) = L(s_0) \cup \{ enabled(\beta) \mid \beta \in Act(s_0) \}$

Action vs State based Fairness 2

Theorem. traces(\mathcal{M}) = traces(\mathcal{M} '). Moreover, strong fairness for a set of actions $F \subseteq A$ can be described by the LTL formula: $strongFair_F \equiv \mathbf{G} \mathbf{F}$ enabled(F) \rightarrow taken(F) [similar for weak fairness and unconditional fairness]

Theorem. Let \mathcal{M} be a transition system without terminal states and let φ be a LTL formula and let \mathcal{F} be a fairness assumption that can be modeled by a LTL formula ψ . Then:

 $\mathcal{M} \vDash_{\mathcal{T}} \varphi \text{ if and only if } \mathcal{M} \vDash \psi \rightarrow \varphi$

Be careful about complexity

Best LTL model checking algorithms *are exponential on the size of the formula* φ to be verified.

If fairness constraints are modeled by complex LTL formula ψ , the **computational cost** to solve the model checking problem

 $\mathcal{M} \vDash \psi \rightarrow \varphi$

could be huge!

Lesson 6c:

Fairness in CTL & CTL Model Checking with fairness constraints

Strong Fairness in CTL

We will consider strong fairness constraints of the form:

$$sfair = \bigwedge_{1 \le i \le k} \mathbf{G} \, \mathbf{F} \, \varphi_i \ \to \mathbf{G} \, \mathbf{F} \, \psi_i$$

Where φ_i and ψ_i are CTL formulas (without fairness). Observe that being CTL formulas, φ_i and ψ_i identify a set of states of a Kripke structure \mathcal{M} : Sat(φ_i) = { $s \mid \mathcal{M}, s \models \varphi_i$ }.

On the other hand, given a path $\pi = s_0 s_1 s_2 \dots$, we have:

$$\pi \vDash_{\mathrm{LTL}} \mathbf{G} \mathbf{F} \varphi_i \to \mathbf{G} \mathbf{F} \psi_i$$

if for all $1 \le i \le k$, there exists j such that $s_j \models_{\text{CTL}} \varphi_i$ for **finitely many** indices j, or $s_j \models_{\text{CTL}} \psi_i$ for **infinitely many** indices (remember that $a \rightarrow b \equiv \neg a \lor b$).

A path π is fair \mathcal{M} , if $\pi \vDash_{\text{LTL}} \mathbf{G} \mathbf{F} \varphi_i \rightarrow \mathbf{G} \mathbf{F} \psi_i$. We denote with:

- fairPaths(s) the set of fair paths starting in a state s,
- fairPaths(M) the set of fair paths starting in an initial state s₀ of M.

Fairness is not expressible in CTL

Formulas of the form **G F** $\phi \rightarrow$ **G F** ψ are not in CTL, because:

- 1. The formula **G F** φ has two consecutive temporal operators
- 2. The boolean connective \rightarrow is applied to two path formulas

In CTL we **must change the semantics** of **E** and **A** stipulating that they quantify over **fair paths**.

We define a new \models_F semantic satisfaction judgement:

M, s ⊨_F p ⇔ there exists a fair path starting from s and p ∈ L(s).
M, s ⊨_F E(g₁) ⇔ there exists a fair path π starting from s such that π ⊨_F g₁.
M, s ⊨_F A(g₁) ⇔ for all fair paths π starting from s, π ⊨_F g₁.

Observe that 1. influence indirectly also the semantics of temporal operators **X** or **U**!!!

CTL model checking with fairness

Theorem. The CTL model checking problem **with fairness** can be reduced to:

- 1. The CTL model checking problem **without fairness**, and
- 2. The problem of computing $Sat_{fair}(\mathbf{E} \mathbf{G} a)$ for some $a \in AP$.

Proof: This approach is quite straightforward for the propositional logic fragment, for example $\text{Sat}_{\text{fair}}(a)$ if $a \in L(s)$ and there exists a fair path starting in *s*, that is $\mathcal{M}, s \vDash_{\text{fair}} \mathbf{E} \mathbf{G}$ true.

Similarly, of course, for $\mathcal{M}, s \vDash_{\text{fair}} \mathbf{E} \mathbf{X} f$: there must be a fair path starting in *s* such that $s_1 \vDash f$ and $\mathcal{M}, s_1 \vDash_{\text{fair}} \mathbf{E} \mathbf{G}$ true.

As for $\mathcal{M}, s \vDash_{\text{fair}} \mathbf{E}[f_1 \mathbf{U} f_2]$ there must be a fair path starting in s such that there exist $\mathcal{M}, s_n \vDash_{\text{fair}} f_2$ and $\mathcal{M}, s_i \vDash_{\text{fair}} f_1 \mathcal{M}$, for all $1 \le i \le n$ and $s_n \vDash_{\text{fair}} \mathbf{E} \mathbf{G}$ true (observe that only the infinite suffix is relevant for fairness).

Obviously, in the iterative CTL algorithm, \mathcal{M} , $s \vDash_{\text{fair}} \mathbf{E} \mathbf{G} f$ is applied when f has been processed, and hence the problem is to check \mathcal{M} , $s \vDash \mathbf{E} \mathbf{G} a_f$ with a_f atomic proposition. \Box

Summing up...

Let a_{fair} be a fresh atomic proposition such that: $a_{\text{fair}} \in L(s)$ if and only if $s \in \text{Sat}_{\text{fair}}(\mathbf{E} \mathbf{G} \text{ true}) \equiv \mathcal{M}, s \vDash_{\text{fair}} \mathbf{E} \mathbf{G}$ true Then:

$$Sat_{fair}(\mathbf{E} \mathbf{X} a) \equiv Sat(\mathbf{E} \mathbf{X} a \land a_{fair})$$
$$Sat_{fair}(\mathbf{E} [a \mathbf{U} a']) \equiv Sat(\mathbf{E} [a \mathbf{U} a' \land a_{fair}])$$

And those on the right-hand side are pure CTL formulas that can be computed by the usual CTL algorithm (see lesson **3**).

Therefore, we are left with the problem of computing $Sat_{fair}(E G a)$

That, in particular, can be used to compute $a_{\text{fair}} \in L(s) \equiv s \in \text{Sat}_{\text{fair}}(\mathbf{E} \mathbf{G} \text{ true})$.

Checking M, $s \models_{fair} EG a$ (1)

Lemma. Let sfair = $\bigwedge_{1 \le i \le k} \mathbf{G} \mathbf{F} a_i \rightarrow \mathbf{G} \mathbf{F} b_i$ be a fair constraint. Then $\mathcal{M}, s \vDash_{\text{sfair}} \mathbf{E} \mathbf{G} a$ if and only if there exists a finite path $s_0 s_1 \dots s_n$ and a cycle $s'_0 s'_1 \dots s'_r$ such that:

- *i.* $s=s_0$ and $s'_0=s'_r$
- *ii.* $s_i \models a$ for all $0 \le i \le n$ and $s'_i \models a$ for all $0 \le j \le r$
- *iii.* For all $0 \le i \le k$ Sat $(a_i) \cap \{s'_0, s'_1, \dots, s'_r\} = \emptyset$ or Sat $(b_i) \cap \{s_0, s_1, \dots, s_n\} \ne \emptyset$

Proof (if): Clearly $s_0s_1...s_n(s'_0s'_1...s'_r)^{\omega}$ is a fair path according to sfair satisfying **EG** *a*.

(**Only if**) \mathcal{M} , $s \models_{\text{sfair}} \text{EG } a$ implies that there exists an infinite fair path $\pi = s_0 s_1 s_2 \dots$ such that $\pi \models_{\text{sfair}} \text{G} a$ and $\pi \models$ sfair. Two cases: 1. $\pi \models \text{G F } a_i$. This implies exists $s' \models b_i$ visited infinitely often in π . Let n and r be the first and second occurrence of s'. Clearly { s_0, s_1, \dots, s_n } and { s_n, s_{n+1}, \dots, s_r } satisfies *iii*.

2. $\pi \nvDash \mathbf{G} \mathbf{F} a_i$. Then there exists *m* such that $s_m, s_{m+1}, \ldots \notin \operatorname{Sat}(b_i)$. There are finitely many states, there is a cycle $s_n, s_{n+1}, \ldots, s_r$ (*n*>*m*) such that $\operatorname{Sat}(a_i) \cap \{s_n, s_{n+1}, \ldots, s_r\} = \emptyset$.

Checking $\mathcal{M}, s \models_{\text{fair}} EG a$ (2)

The previous Lemma can be used as follows. Consider the graph G_a whose nodes $V_a = \{s \mid \mathcal{M}, s \models a\}$ and edges $E_a = \{(s, s') \in R \mid s \in V_a, s' \in V_a\}$.

Each infinite path in G_a is a path in \mathcal{M} satisfying **G** *a*. Conversely, each path in \mathcal{M} satisfying **G** *a* is a path in G_a .

 \mathcal{M} , $s \models_{\text{sfair}} \mathbf{EG} a$ if and only if there exists a nontrivial SCC *C* in G_a reachable from *s* and a set of nodes $D \subseteq C$ such that for all $0 \le i \le k$, $D \cap \text{Sat}(a_i) = \emptyset$ or $D \cap \text{Sat}(b_i) \neq \emptyset$.

 $Sat_{fair}(EG a) = \{s \mid exists C reachable from s in G_a\}$

Unconditional Fairness: in this case, *a_i* is true for all *i*. Observe that in this case, fair path **corresponds to accepting runs of a Generalised Büchi automaton**.

Example: unconditional fairness



 G_1 satisfy unconditional fairness constraint **G F** $b_1 \wedge \mathbf{G} \mathbf{F} b_2$ because there is the SCC { s_2 , s_3 , s_4 }.

By contrast, G_2 does not satisfy **G F** $b_1 \wedge \mathbf{G} \mathbf{F} b_2$ because there is the SCC { s_2, s_3, s_4 } that contains b_2 and the SCC { s_1 } that contains only b_1 , but **no one of them contains both** b_1 and b_2 .

Lesson 6

That's all Folks...

... Questions?