Timed Automata

Lecture #16 of Advanced Model Checking

Joost-Pieter Katoen

Lehrstuhl 2: Software Modeling & Verification

E-mail: katoen@cs.rwth-aachen.de

January 8, 2006

Clock constraints

- *Clock constraints* over set C of clocks are defined by:

$$g ::= \text{true} \mid x < c \mid x - y < c \mid x \leqslant c \mid x - y \leqslant c \mid \neg g \mid g \wedge g$$

  - where $c \in \mathbb{N}$ and clocks $x, y \in C$
  - rational constants would do; neither reals nor addition of clocks!
  - let $CC(C)$ denote the set of clock constraints over C
  - shorthands: $x \geqslant c$ denotes $\neg(x < c)$ and $x \in [c_1, c_2)$ or $c_1 \leqslant x < c_2$ denotes $\neg(x < c_1) \wedge (x < c_2)$

- *Atomic clock constraints* do not contain true, $\neg$ and $\wedge$

  - let $ACC(C)$ denote the set of atomic clock constraints over C

Timed automaton

A *timed automaton* is a tuple

$$TA = \left(Loc, Act, C, \rightsquigarrow, Loc_0, Inv, AP, L\right) \quad \text{where:}$$

- Loc is a finite set of *locations*
- $Loc_0 \subseteq Loc$ is a set of initial locations
- C is a finite set of *clocks*
- $L : Loc \to 2^{AP}$ is a labeling function for the locations
- $\rightsquigarrow \subseteq Loc \times CC(C) \times Act \times 2^C \times Loc$ is a transition relation, and
- $Inv : Loc \to CC(C)$ is an *invariant*-assignment function

Intuitive interpretation

- Edge $\ell \xrightarrow{g:\alpha,C'} \ell'$ means:
  - action α is*enabled* once guard g holds
  - when moving from location ℓ to ℓ', any clock in C' will be *reset* to zero
- *Inv*(ℓ) constrains the amount of time that may be spent in location ℓ
  - once the invariant *Inv*(ℓ) becomes invalid, the location ℓ must be left immediately
  - if this is not possible – no enabled outgoing transition – no further progress is possible

Example: the gate

Clock valuations

- A *clock valuation* η for set C of clocks is a function $\eta : C \longrightarrow \mathbb{R}_{\geqslant 0}$
  - assigns to each clock $x \in C$ its current value $\eta(x)$
- Clock valuation $\eta{+}d$ for $d \in \mathbb{R}_{\geqslant 0}$ is defined by:
  - $(\eta{+}d)(x) = \eta(x) + d$ for all clocks $x \in C$
- Clock valuation reset x in η for clock x is defined by:

$$(\text{reset } x \text{ in } \eta)(y) = \begin{cases} \eta(y) & \text{if } y \neq x \\ 0 & \text{if } y = x \end{cases}$$

  - reset x in (reset y in η) is abbreviated by reset x, y in η

Semantics of clock constraints

Let $\models \subseteq \mathit{Eval}(C) \times \mathit{CC}(C)$ be defined by:

$$
\begin{array}{lll}
\eta \models \text{true} & & \\
\eta \models x < c & \text{iff} & \eta(x) < c \\
\eta \models x \leqslant c & \text{iff} & \eta(x) \leqslant c \\
\eta \models x - y < c & \text{iff} & \eta(x) - \eta(y) < c \\
\eta \models x - y \leqslant c & \text{iff} & \eta(x) - \eta(y) \leqslant c \\
\eta \models \neg\, g & \text{iff} & \eta \not\models g \\
\eta \models g \wedge g' & \text{iff} & \eta \models g \ \wedge \ \eta \models g'
\end{array}
$$

Semantics

For timed automaton $\mathit{TA} = (\mathit{Loc}, \mathit{Act}, C, \hookrightarrow, \mathit{Loc}_0, \mathit{Inv}, \mathit{AP}, L)$:

Transition system $\mathit{TS}(\mathit{TA}) = (S, \mathit{Act}', \rightarrow, I, \mathit{AP}', L')$ where:

- $S = \mathit{Loc} \times \mathit{val}(C)$, state $s = \langle \ell, \eta \rangle$ for location ℓ and clock valuation η
- $\mathit{Act}' = \mathit{Act} \cup \mathbb{R}_{\geqslant 0}$, (discrete) actions and time passage actions
- $I = \{ \langle \ell_0, \eta_0 \rangle \mid \ell_0 \in \mathit{Loc}_0 \ \wedge \ \eta_0(x) = 0$ for all $x \in C \}$
- $\mathit{AP}' = \mathit{AP} \cup \mathit{ACC}(C)$
- $L'(\langle \ell, \eta \rangle) = L(\ell) \cup \{ g \in \mathit{ACC}(C) \mid \eta \models g \}$
- $\rightarrow$ is the transition relation defined on the next slide

Semantics

The transition relation $\rightarrow$ is defined by the following two rules:

- Discrete transition: $\langle \ell, \eta \rangle \xrightarrow{d} \langle \ell', \eta' \rangle$ if all following conditions hold:
  - there is an edge labeled $(g : \alpha, D)$ from location ℓ to ℓ' such that:
  - g is satisfied by η, i.e., $\eta \models g$
  - $\eta' = \eta$ with all clocks in D reset to 0, i.e., $\eta' =$ reset D in η
  - η' fulfills the invariant of location ℓ', i.e., $\eta' \models \mathit{Inv}(\ell')$
- Delay transition: $\langle \ell, \eta \rangle \xrightarrow{\alpha} \langle \ell, \eta+d \rangle$ for positive real d
  - if for any $0 \leqslant d' \leqslant d$ the invariant of ℓ holds for $\eta+d'$, i.e. $\eta+d' \models \mathit{Inv}(\ell)$

Example

Time divergence

- Let for any $t < d$, for fixed $d \in \mathbb{R}_{>0}$, clock valuation $\eta{+}t \models \mathit{Inv}(\ell)$
- A possible execution fragment starting from the location ℓ is:

$$\langle \ell, \eta \rangle \xrightarrow{d_1} \langle \ell, \eta{+}d_1 \rangle \xrightarrow{d_2} \langle \ell, \eta{+}d_1{+}d_2 \rangle \xrightarrow{d_3} \langle \ell, \eta{+}d_1{+}d_2{+}d_3 \rangle \xrightarrow{d_4} \ldots$$

  - where $d_i > 0$ and the infinite sequence $d_1 + d_2 + \ldots$ *converges* towards d
  - such path fragments are called *time-convergent*

  $\Rightarrow$ time advances only up to a certain value

- Time-convergent execution fragments are unrealistic and *ignored*
  - much like unfair paths (as we will see later on)

Time divergence

- Infinite path fragment π is *time-divergent* if $ExecTime(\pi) = \infty$;
- The function $ExecTime : \mathit{Act} \cup \mathbb{R}_{>0} \to \mathbb{R}_{\geqslant 0}$ is defined as:

$$ExecTime(\tau) = \begin{cases} 0 & \text{if } \tau \in \mathit{Act} \\ d & \text{if } \tau = d \in \mathbb{R}_{>0} \end{cases}$$

- For infinite execution fragment $\rho = s_0 \xrightarrow{\tau_1} s_1 \xrightarrow{\tau_2} s_2 \ldots$ in $\mathit{TS}(\mathit{TA})$ let:

$$ExecTime(\rho) = \sum_{i=0}^{\infty} ExecTime(\tau_i)$$

  - for path fragment π in $\mathit{TS}(\mathit{TA})$ induced by ρ: $ExecTime(\pi) = ExecTime(\rho)$

- For state s in $\mathit{TS}(\mathit{TA})$: $\mathit{Paths}_{div}(s) = \{\, \pi \in \mathit{Paths}(s) \mid \pi \text{ is time-divergent} \,\}$

Example: light switch

The path π in *TS*$(Switch)$ in which on- and of-periods of one minute alternate:

$$\pi = \langle off, 0\rangle \langle off, 1\rangle \langle on, 0\rangle \langle on, 1\rangle \langle off, 1\rangle \langle off, 2\rangle \langle on, 0\rangle \langle on, 1\rangle \langle off, 2\rangle \ldots$$

is *time-divergent* as $ExecTime(\pi) = 1 + 1 + 1 + \ldots = \infty$.

The path:

$$\pi' = \langle off, 0\rangle \langle off, 1/2\rangle \langle off, 3/4\rangle \langle off, 7/8\rangle \langle off, 15/16\rangle \ldots$$

is *time-convergent*, since $ExecTime(\pi') = \sum_{i \geqslant 1} \left(\frac{1}{2}\right)^i = 1 < \infty$

Timelock

- State $s \in \mathit{TS}(\mathit{TA})$ contains a *timelock* if $\mathit{Paths}_{div}(s) = \varnothing$
  - there is no behavior in s where time can progress *ad infinitum*
  - clearly: any terminal state contains a timelock (but also non-terminal states may do)
- *TA* is *timelock-free* if no state in $\mathit{Reach}(\mathit{TS}(\mathit{TA}))$ contains a timelock
- Timelocks are considered as *modeling flaws* that should be avoided
  - like deadlocks, we need mechanisms to check their presence

Example

Zenoness

- A *TA* that performs infinitely many actions in finite time is *Zeno*
- Path π in $TS(TA)$ is *Zeno* if:
  - it is time-convergent, and
  - infinitely many actions $\alpha \in Act$ are executed along π
- *TA* is *non-Zeno* if there does not exist an initial Zeno path in $TS(TA)$
  - any π in $TS(TA)$ is time-divergent or
  - is time-convergent with nearly all (i.e., all except for finitely many) transitions being delay transitions
- Zeno paths are considered as *modeling flaws* that should be avoided
  - like timelocks (and deadlocks), we need mechanisms to check Zenoness

Example

A sufficient criterion for Zenoness

Let *TA* with set C of clocks such that for every control cycle:

$$\ell_0 \overset{g_1:\alpha_1,C_1}{\rightsquigarrow} \ell_1 \overset{g_2:\alpha_2,C_2}{\rightsquigarrow} \ldots \overset{g_n:\alpha_n,C_n}{\rightsquigarrow} \ell_n$$

there exists a clock $x \in C$ such that:

1. $x \in C_i$ for some $0 < i \leqslant n$, and

2. for all clock evaluations η:

$$\eta(x) < 1 \text{ implies } (\eta \not\models g_j \text{ or } \eta \not\models \mathit{Inv}(\ell_j)), \text{ for some } 0 < j \leqslant n$$

Then: *TA* is *non-Zeno*

Proof

Example

Timelock, time-divergence and Zenoness

- A timed automaton is adequately modeling a time-critical system whenever it is:

  non-Zeno and timelock-free

- Time-divergent paths will be explicitly excluded for analysis purposes

Timed CTL

Syntax of TCTL *state-formulas* over *AP* and set C:

$$\Phi ::= \text{true} \mid a \mid g \mid \Phi \wedge \Phi \mid \neg\Phi \mid \exists\varphi \mid \forall\varphi$$

where $a \in AP$, $g \in ACC(C)$ and φ is a path-formula defined by:

$$\varphi ::= \Diamond^{J}\Phi$$

where $J \subseteq \mathbb{R}_{\geqslant 0}$ is an interval whose bounds are naturals

$\Diamond^{J}\Phi$ asserts that a Φ-state is reached at time instant $t \in J$

Forms of J: $[n, m]$, $(n, m]$, $[n, m)$ or (n, m) for $n, m \in \mathbb{N}$ and $n \leqslant m$

for right-open intervals, $m = \infty$ is also allowed

Some abbreviations

"Always" is obtained in the following way:

$$\exists \Box^J \Phi \;=\; \neg \forall \Diamond^J \neg \Phi \quad \text{and} \quad \forall \Box^J \Phi \;=\; \neg \exists \Diamond^J \neg \Phi$$

$\exists \Box^J \Phi$ asserts that for some path during the interval J, Φ holds

$\forall \Box^J \Phi$ requires this to hold for all paths

Standard until-operator is obtained as follows:

$$\Diamond \Phi = \Diamond^{[0,\infty)} \Phi \quad \text{and} \quad \Box \Phi = \Box^{[0,\infty)} \Phi$$

Timed properties in TCTL

Semantics of TCTL

For state $s = \langle \ell, \eta \rangle$ in *TS*(*TA*) the satisfaction relation $\models$ is defined by:

$$
\begin{array}{lll}
s \models \text{true} & & \\
s \models a & \text{iff} & a \in L(\ell) \\
s \models g & \text{iff} & \eta \models g \\
s \models \neg \Phi & \text{iff} & \text{not } s \models \Phi \\
s \models \Phi \wedge \Psi & \text{iff} & (s \models \Phi) \text{ and } (s \models \Psi) \\
s \models \exists \varphi & \text{iff} & \pi \models \varphi \text{ for some } \pi \in \mathit{Paths}_{div}(s) \\
s \models \forall \varphi & \text{iff} & \pi \models \varphi \text{ for all } \pi \in \mathit{Paths}_{div}(s)
\end{array}
$$

path quantification over time-divergent paths only

The $\Longrightarrow$ relation

For infinite path fragments in $TS(TA)$ performing ∞ many actions let:

$$s_0 \overset{d_0}{\Longrightarrow} s_1 \overset{d_1}{\Longrightarrow} s_2 \overset{d_2}{\Longrightarrow} \ldots \qquad \text{with } d_0, d_1, d_2 \ldots \geqslant 0$$

denote the equivalence class containing all infinite path fragments induced by execution fragments of the form:

$$s_0 \underbrace{\overset{d_0^1}{\rightarrow} \ldots \overset{d_0^{k_0}}{\rightarrow}}_{\substack{\text{time passage of} \\ d_0 \text{ time-units}}} s_0{+}d_0 \overset{\alpha_1}{\longrightarrow} s_1 \underbrace{\overset{d_1^1}{\rightarrow} \ldots \overset{d_1^{k_1}}{\rightarrow}}_{\substack{\text{time passage of} \\ d_1 \text{ time-units}}} s_1{+}d_1 \overset{\alpha_2}{\longrightarrow} s_2 \underbrace{\overset{d_2^1}{\rightarrow} \ldots \overset{d_2^{k_2}}{\rightarrow}}_{\substack{\text{time passage of} \\ d_2 \text{ time-units}}} s_2{+}d_2 \overset{\alpha_3}{\longrightarrow} \ldots$$

where $k_i \in \mathbb{N}$, $d_i \in \mathbb{R}_{\geqslant 0}$ and $\alpha_i \in \mathit{Act}$ such that $\sum_{j=1}^{k_i} d_i^j = d_i$.

For $\pi \in s_0 \overset{d_0}{\Longrightarrow} s_1 \overset{d_1}{\Longrightarrow} \ldots$ we have $ExecTime(\pi) = \sum_{i \geqslant 0} d_i$

Semantics of TCTL

For time-divergent path $\pi \in s_0 \xRightarrow{d_0} s_1 \xRightarrow{d_1} \ldots$:

$\pi \models \Diamond^{J} \Phi$

iff

$\exists\, i \geqslant 0.\ s_i{+}d \models \Phi$ for some $d \in [0, d_i]$ with $\sum_{j=0}^{i-1} d_j + d \in J$

where for $s_i = \langle \ell_i, \eta_i \rangle$ we have $s_i{+}d = \langle \ell_i, \eta_i{+}d \rangle$

TCTL-semantics for timed automata

- Let *TA* be a timed automaton with clocks C and locations *Loc*
- For TCTL-state-formula Φ, the *satisfaction set* *Sat*(Φ) is defined by:

$$\mathit{Sat}(\Phi) \;=\; \{\, s \in \mathit{Loc} \times \mathit{Eval}(C) \mid s \models \Phi \,\}$$

- *TA* satisfies TCTL-formula Φ iff Φ holds in all initial states of *TA*:

$$\mathit{TA} \models \Phi \quad \text{if and only if} \quad \forall \ell_0 \in \mathit{Loc}_0.\, \langle \ell_0, \eta_0 \rangle \models \Phi$$

where $\eta_0(x) = 0$ for all $x \in C$

Example

Timed CTL versus CTL

- Due to ignoring time-convergent paths in TCTL semantics possibly:

$$\underbrace{TS(TA) \models_{\text{TCTL}} \forall\varphi}_{\text{TCTL semantics}} \quad \text{but} \quad \underbrace{TS(TA) \not\models_{\text{CTL}} \forall\varphi}_{\text{CTL semantics}}$$

  - CTL semantics considers all paths, timed CTL only time-divergent paths

- For $\Phi = \forall\Box\big(on \longrightarrow \forall\Diamond off\big)$ and the light switch

$$TS(Switch) \models_{\text{TCTL}} \Phi \quad \text{whereas} \quad TS(TA) \not\models_{\text{CTL}} \Phi$$

  - there are time-convergent paths on which location *on* is never left

Characterizing timelock

- TCTL semantics is also well-defined for *TA* with timelock
- A state is *timelock-free* if and only if it satisfies $\exists\Box$true
  - some time-divergent path satisfies $\Box$true, i.e., there is $\geqslant 1$ time-divergent path
  - note: for fair CTL, the states in which a fair path starts also satisfy $\exists\Box$true
- *TA* is timelock-free iff $\forall s \in \mathit{Reach}(\mathit{TS}(\mathit{TA}))$: $s \models \exists\Box$true
- Timelocks can thus be checked by model checking!