Playing optimally on Timed Automata with Random Delays

Nathalie Bertrand and Sven Schewe

Inria Rennes - University of Liverpool

Formats 2012

Stochastic timed automata [BBBBG-fsttcs07&lics08]

Timed automata with random delays and random actions.

Semantics: from any state

1. sample a delay, according to a fixed probability distribution
2. select randomly an enabled edge

   $\rightarrow$ Infinite-state Markov chain.

Stochastic timed automata [BBBBG-fsttcs07&lics08]

Timed automata with random delays and random actions.

Semantics: from any state

1. sample a delay, according to a fixed probability distribution
2. select randomly an enabled edge

$\rightarrow$ Infinite-state Markov chain.

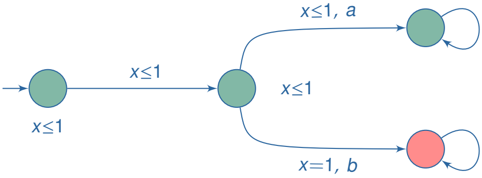

uniform distributions on delays and actions

Stochastic timed automata [BBBBG-fsttcs07&lics08]

Timed automata with random delays and random actions.

Semantics: from any state

1. sample a delay, according to a fixed probability distribution
2. select randomly an enabled edge

   $\rightarrow$ Infinite-state Markov chain.

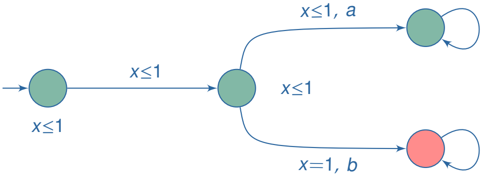

uniform distributions on delays and actions

almost-surely "**G** green" is satisfied

Almost-sure model checking of STA

Almost-sure model checking

Is φ almost-surely is satisfied?

Almost-sure model checking of STA

Almost-sure model checking

Is φ almost-surely is satisfied?

Decidability of almost-sure model-checking

One can decide whether $\mathbb{P}(\varphi) = 1$ if

- safety properties, or
- 1-clock timed automaton, or
- reactive timed automata

Almost-sure model checking of STA

Almost-sure model checking

Is φ almost-surely is satisfied?

Decidability of almost-sure model-checking

One can decide whether $\mathbb{P}(\varphi) = 1$ if

- safety properties, or
- 1-clock timed automaton, or
- reactive timed automata

Technique: build an equivalent discrete finite-state Markov chain.

Almost-sure model checking of STA

Almost-sure model checking

Is φ almost-surely is satisfied?

Decidability of almost-sure model-checking

One can decide whether $\mathbb{P}(\varphi) = 1$ if

- safety properties, or
- 1-clock timed automaton, or
- reactive timed automata

Technique: build an equivalent discrete finite-state Markov chain.

Extends to quantitative properties for 1-clock timed automata under some additional assumptions.

STA as modelling formalism

IP address allocation protocol

device entering a network

- selects an IP at random
- sends message to find out whether the IP is already taken

STA as modelling formalism

IP address allocation protocol

device entering a network

- selects an IP at random
- sends message to find out whether the IP is already taken

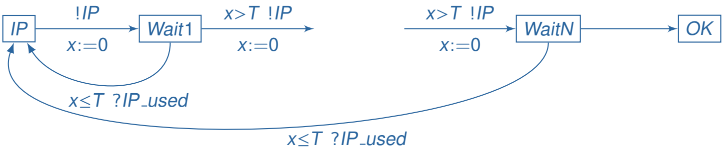

Timed automata Markov decision processes

From any state,

1. sample a delay, according to a fixed probability distribution
2. select an enabled action

Timed automata Markov decision processes

From any state,

1. sample a delay, according to a fixed probability distribution
2. select an enabled action

Assumptions

- reactive timed automaton
  for every state and every delay,
  some action is enabled
- exponential distributions on delays

Timed automata Markov decision processes

From any state,

1. sample a delay, according to a fixed probability distribution
2. select an enabled action

Assumptions

- reactive timed automaton
  for every state and every delay,
  some action is enabled
- exponential distributions on delays

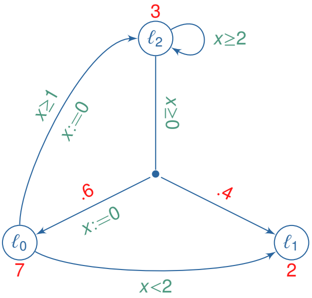

Timed automata Markov decision processes

From any state,

1. sample a delay, according to a fixed probability distribution
2. select an enabled action

Assumptions

- reactive timed automaton
  for every state and every delay,
  some action is enabled
- exponential distributions on delays

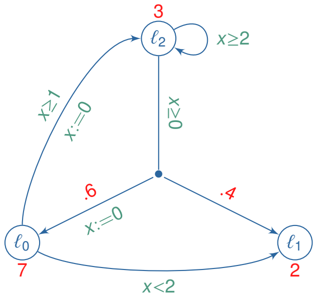

Special cases

CTMDP = TAMDP without clocks
STA = TAMDP with a single action

Meaningful schedulers

Scheduler

- resolves non-determinism
- associates with each prefix run and delay, a distribution over enabled actions

$$\sigma : \mathsf{Runs}(\mathcal{M}) \times \mathbb{R}_{\geq 0} \to \mathsf{Dist}(\mathit{Act})$$

Late scheduler: decision is made right before discrete transition

Meaningful schedulers

Scheduler

- resolves non-determinism
- associates with each prefix run and delay, a distribution over enabled actions

$$\sigma : \mathrm{Runs}(\mathcal{M}) \times \mathbb{R}_{\geq 0} \to \mathrm{Dist}(\mathit{Act})$$

Late scheduler: decision is made right before discrete transition

Given TAMDP $\mathcal{M}$, σ defines a probability measure $\mathbb{P}_\sigma$ over runs.

Measurability constraints on σ to be well-defined

Problem statement

Paths reaching goal G within T, from (ℓ_0, v_0)

$$\mathsf{Reach}_{\mathcal{M}}(\ell, v, G, T) = \{(\ell_0, v_0) \xrightarrow{t_0, e_0, p_0} (\ell_1, v_1) \cdots (\ell_n, v_n) \mid \exists i \leq n,\ \ell_i \in G \text{ and } \sum_{j<i} t_j \leq T\}.$$

Optimal probability

$$\mathsf{Opt}_{\mathcal{M}}(\ell, v, G, T) = \sup_{\sigma}\ \mathbb{P}_{\sigma}(\mathsf{Reach}_{\mathcal{M}}(\ell, v, G, T)).$$

Problem statement

Paths reaching goal G within T, from (ℓ_0, v_0)

$$\mathrm{Reach}_{\mathcal{M}}(\ell, v, G, T) = \{(\ell_0, v_0) \xrightarrow{t_0, e_0, p_0} (\ell_1, v_1) \cdots (\ell_n, v_n) \mid \exists i \leq n,\ \ell_i \in G \text{ and } \sum_{j<i} t_j \leq T\}.$$

Optimal probability

$$\mathrm{Opt}_{\mathcal{M}}(\ell, v, G, T) = \sup_{\sigma} \ \mathbb{P}_{\sigma}(\mathrm{Reach}_{\mathcal{M}}(\ell, v, G, T)).$$

Fundamental questions

- Is the optimal realized? $\sup = \max$?
- For what class of schedulers?

Approximations from below

Paths reaching G within T and in less than N discrete steps
$\text{Reach}^N(\ell, v, G, T)$

$$\text{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = \sup_{\sigma} \mathbb{P}_\sigma(\text{Reach}^N_{\mathcal{M}}(\ell, v, G, T)).$$

Approximations from below

Paths reaching G within T and in less than N discrete steps
$\text{Reach}^N(\ell, v, G, T)$

$$\text{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = \sup_{\sigma} \mathbb{P}_\sigma(\text{Reach}^N_{\mathcal{M}}(\ell, v, G, T)).$$

Optimal probability is the limit of step-bounded optimals

$$\lim_{N\to\infty} \text{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = \text{Opt}_{\mathcal{M}}(\ell, v, G, T)$$

Approximations from below

Paths reaching G within T and in less than N discrete steps
$\mathrm{Reach}^N(\ell, v, G, T)$

$$\mathrm{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = \sup_\sigma \mathbb{P}_\sigma(\mathrm{Reach}^N_{\mathcal{M}}(\ell, v, G, T)).$$

Optimal probability is the limit of step-bounded optimals

$$\lim_{N\to\infty} \mathrm{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = \mathrm{Opt}_{\mathcal{M}}(\ell, v, G, T)$$

Sketch

- fix scheduler σ ε-optimal

$$\mathbb{P}_\sigma(\mathrm{Reach}_{\mathcal{M}}) \geq \mathrm{Opt}_{\mathcal{M}} - \varepsilon$$

- fix step-bound N s.t.

$$\mathbb{P}_\sigma(\mathrm{Reach}^N_{\mathcal{M}}) \geq \mathbb{P}_\sigma(\mathrm{Reach}_{\mathcal{M}}) - \varepsilon$$

Properties

Characterization

$$\mathrm{Opt}^0_{\mathcal{M}}(\ell, v, G, T) = 0 \text{ if } \ell \notin G$$

$$\mathrm{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = 1 \text{ if } \ell \in G$$

$$\mathrm{Opt}^{N+1}_{\mathcal{M}}(\ell, v, G, T) = \int_0^T \max_{e \in E} \sum_{(\ell,v) \xrightarrow{t,e,p} (\ell',v')} p \cdot \mathrm{Opt}^N_{\mathcal{M}}(\ell', v', G, T-t) \cdot \Lambda(\ell) \cdot e^{-\Lambda(\ell)t} \mathrm{d}t$$

Properties

Characterization

$$\mathrm{Opt}^0_{\mathcal{M}}(\ell, v, G, T) = 0 \text{ if } \ell \notin G$$

$$\mathrm{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = 1 \text{ if } \ell \in G$$

$$\mathrm{Opt}^{N+1}_{\mathcal{M}}(\ell, v, G, T) = \int_0^T \max_{e \in E} \sum_{(\ell,v) \xrightarrow{t,e,p} (\ell',v')} p \cdot \mathrm{Opt}^N_{\mathcal{M}}(\ell', v', G, T-t) \cdot \Lambda(\ell) \cdot e^{-\Lambda(\ell)t} \mathrm{d}t$$

Uniform continuity

$\mathrm{Opt}_{\mathcal{M}}(\ell, v+t, G, T-t)$ is uniformly continuous in t and v.

Properties

Characterization

$$\mathrm{Opt}^0_{\mathcal{M}}(\ell, v, G, T) = 0 \text{ if } \ell \notin G$$

$$\mathrm{Opt}^N_{\mathcal{M}}(\ell, v, G, T) = 1 \text{ if } \ell \in G$$

$$\mathrm{Opt}^{N+1}_{\mathcal{M}}(\ell, v, G, T) = \int_0^T \max_{e \in E} \sum_{(\ell,v) \xrightarrow{t,e,p} (\ell',v')} p \cdot \mathrm{Opt}^N_{\mathcal{M}}(\ell', v', G, T - t) \cdot \Lambda(\ell) \cdot e^{-\Lambda(\ell)t} \mathrm{d}t$$

Uniform continuity

$\mathrm{Opt}_{\mathcal{M}}(\ell, v + t, G, T - t)$ is uniformly continuous in t and v.

Key idea for uniform continuity in v

- $\mathrm{Opt}^{N+1}_{\mathcal{M}}(\ell, v) - \mathrm{Opt}^{N+1}_{\mathcal{M}}(\ell, w) = \int \max_{(\ell,v) \xrightarrow{t,e}} \mathrm{Opt}^N_{\mathcal{M}} \cdots - \int \max_{(\ell,w) \xrightarrow{t,e}} \mathrm{Opt}^N_{\mathcal{M}} \cdots$
- bound the measure μ of delays for which enabled edges differ

$$\|v - w\| < \delta \quad \Rightarrow \quad m < n\lceil T \rceil \delta$$

Existence of optimal schedulers

Optimal schedulers exist

For every TAMDP $\mathcal{M}$, reachability goal G and time-bound T, there exists a measurable scheduler σ such that:

$$\mathbb{P}_\sigma(\mathsf{Reach}_{\mathcal{M}}(\ell_0, 0^X, G, T)) = \mathsf{Opt}_{\mathcal{M}}(\ell_0, 0^X, G, T).$$

Existence of optimal schedulers

Optimal schedulers exist

For every TAMDP $\mathcal{M}$, reachability goal G and time-bound T, there exists a measurable scheduler σ such that:

$$\mathbb{P}_\sigma(\mathsf{Reach}_{\mathcal{M}}(\ell_0, 0^X, G, T)) = \mathsf{Opt}_{\mathcal{M}}(\ell_0, 0^X, G, T).$$

Illustration of the proof

- D_a area in T^{n+1} where action a is optimal
- by continuity, a is also optimal on $\bar{D}_a$
- arbitrary order on actions as tie-breaker

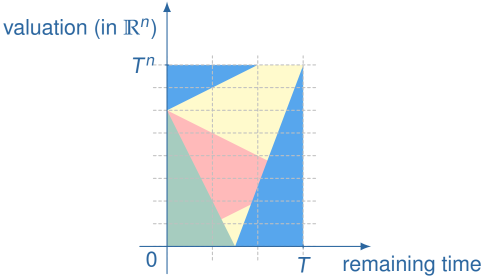

Existence of optimal schedulers

Optimal schedulers exist

For every TAMDP $\mathcal{M}$, reachability goal G and time-bound T, there exists a measurable scheduler σ such that:

$$\mathbb{P}_\sigma(\mathrm{Reach}_{\mathcal{M}}(\ell_0, 0^X, G, T)) = \mathrm{Opt}_{\mathcal{M}}(\ell_0, 0^X, G, T).$$

Illustration of the proof

- D_a area in T^{n+1} where action a is optimal
- by continuity, a is also optimal on $\bar{D}_a$
- arbitrary order on actions as tie-breaker

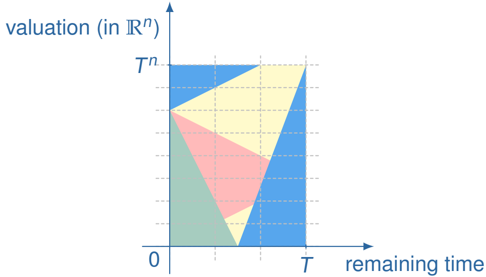

- non-constructive existence proof
- provides a memoryless deterministic measurable scheduler

Time unbounded reachability

Optimal schedulers may not exist for time unbounded reachability.

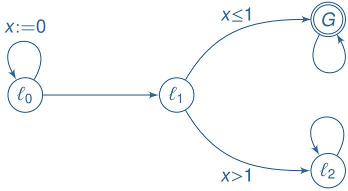

Time unbounded reachability

Optimal schedulers may not exist for time unbounded reachability.

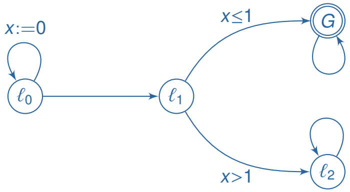

- In ℓ_1, the smaller x the greater the probability.

Time unbounded reachability

Optimal schedulers may not exist for time unbounded reachability.

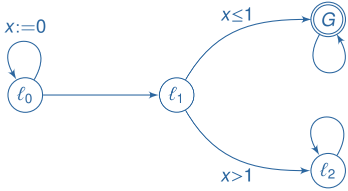

- In ℓ_1, the smaller x the greater the probability.
- In ℓ_0, for any sampled delay $t > 0$, a smaller delay can eventually be obtained by looping on ℓ_0.

Simpler schedulers

Optimal polyhedral schedulers may not exist.

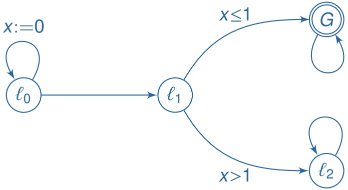

constant rate $\Lambda = 1$
time-bound $T = 1$

Simpler schedulers

Optimal polyhedral schedulers may not exist.

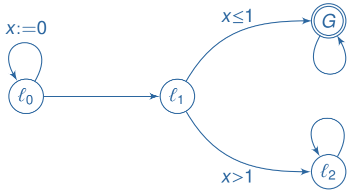

constant rate $\Lambda = 1$
time-bound $T = 1$

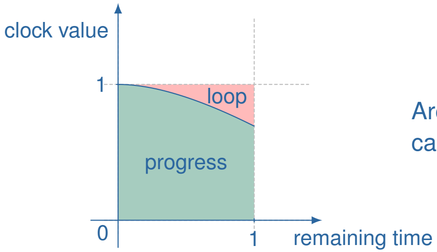

Areas where each action is optimal cannot be described by polyhedra.

Conclusion

- Timed automata MDP model
  - random delays and nondeterministic actions
  - encompass CTMDPs and stochastic timed automata

Conclusion

- Timed automata MDP model
  - random delays and nondeterministic actions
  - encompass CTMDPs and stochastic timed automata

- Other real-time models with probabilities and non-determinism
  - Stochastic timed games [BF-icalp09]
    turn-based game between 2 players (choosing delays and actions) in a randomized environment
  - Stochastic real-time games [BKKKR-concur10]
    CTMDPs with an objective given by a deterministic timed automaton
  - Markovian timed automata [CHKM-lics11&rp11&cdc11]
    Timed automata with exponentially distributed sojourn time

Conclusion

Contribution on TAMDP model

- existence of optimal schedulers for time-bounded reachability
- extends to 2-player games
- does not extend to time unbounded reachability
- optimal schedulers are not region-based

Conclusion

Contribution on TAMDP model

- existence of optimal schedulers for time-bounded reachability
- extends to 2-player games
- does not extend to time unbounded reachability
- optimal schedulers are not region-based

Open questions

- finiteness of partitionning
- decidability of existence of optimal schedulers for time-unbounded reachability
- subclasses with effective schedulers