Delayed Nondeterminism in Continuous-Time Markov Decision Processes

Martin R. Neuhäußer[1,2] Mariëlle Stoelinga[2] Joost-Pieter Katoen[1,2]

[1]RWTH Aachen University, Germany
[2]University of Twente, The Netherlands

FOSSACS 2009, York, United Kingdom

Continuous-Time Markov Decision Processes: An Example

Imagine you have to come home by 6 pm.

- On your way, you stop at a red traffic light.
- When it turns green, you have two choices:

- Best strategy to meet your family's deadline?

Continuous-Time Markov Decision Processes: An Example

Imagine you have to come home by 6 pm.

- On your way, you stop at a red traffic light.
- When it turns green, you have **two choices**:
  - turn left: $1min$; traffic jam probability $\frac{1}{2}$.
  - turn right: $5min$; traffic jam probability $\frac{1}{9}$.
  - Expected delay in a traffic jam: $30min$.
- Best strategy to meet your family's deadline?

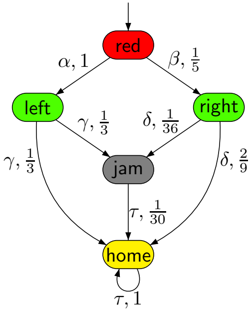

Continuous-Time Markov Decision Processes: An Example

Imagine you have to come home by 6 pm.

- On your way, you stop at a red traffic light.
- When it turns green, you have **two choices**:
  - turn left: $1min$; traffic jam probability $\frac{1}{2}$.
  - turn right: $5min$; traffic jam probability $\frac{1}{9}$.
  - Expected delay in a traffic jam: $30min$.
- Best strategy to meet your family's deadline?

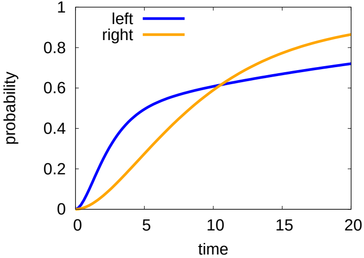

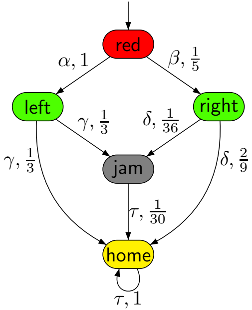

Aim: Maximize the probability to come home in t time units.

Why Continuous-Time Markov Decision Processes?

1. CTMDPs are an important model in
   - stochastic control theory [Qiu et al.]
   - stochastic scheduling [Feinberg et al., Puterman]
2. CTMDPs provide the semantic basis for
   - non-well-specified stochastic activity networks [Sanders et al.]
   - generalised stochastic Petri nets with confusion [Chiola et al.]
   - Markovian process algebras [Hermanns et al., Hillston et al.]

In this talk:

1. Introduction of CTMDPs.
2. Schedulers that resolve the nondeterminism.
3. Probability measures.
4. Delaying nondeterminism.
5. Results and future work.

Why Continuous-Time Markov Decision Processes?

1. CTMDPs are an important model in
   - stochastic control theory [Qiu et al.]
   - stochastic scheduling [Feinberg et al., Puterman]
2. CTMDPs provide the semantic basis for
   - non-well-specified stochastic activity networks [Sanders et al.]
   - generalised stochastic Petri nets with confusion [Chiola et al.]
   - Markovian process algebras [Hermanns et al., Hillston et al.]

In this talk:

1. **Introduction** of CTMDPs.
2. **Schedulers** that resolve the nondeterminism.
3. Probability **measures**.
4. **Delaying nondeterminism**.
5. **Results** and future work.

Continuous Time Markov Decision Process

A tuple $(\mathcal{S}, Act, \mathbf{R}, \nu)$ is a CTMDP if $\mathcal{S}$ is a finite set of **states** and

- $Act = \{\alpha, \beta, \gamma, \dots\}$ is a finite set of **actions** and
- $\mathbf{R} : \mathcal{S} \times Act \times \mathcal{S} \to \mathbb{R}_{\geq 0}$ is a transition **rate matrix** such that
  - $\mathbf{R}(s, \alpha, s') = \lambda$ is the rate of a negative exponential distribution

$$f_X(t) = \begin{cases} \lambda \cdot e^{-\lambda \cdot t} & \text{if } t \geq 0 \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad E[X] = \frac{1}{\lambda}$$

  such that $Act(s) = \{\alpha \in Act \mid \exists s' \in \mathcal{S}.\ \mathbf{R}(s, \alpha, s') > 0\} \neq \emptyset$ for all $s \in \mathcal{S}$.
  - $E(s, \alpha) = \sum_{s' \in \mathcal{S}} \mathbf{R}(s, \alpha, s')$ is the **exit rate** of s under α.

Continuous Time Markov Decision Process

A tuple $(\mathcal{S}, Act, \mathbf{R}, \nu)$ is a CTMDP if $\mathcal{S}$ is a finite set of **states** and

- $Act = \{\alpha, \beta, \gamma, \dots\}$ is a finite set of **actions** and
- $\mathbf{R} : \mathcal{S} \times Act \times \mathcal{S} \to \mathbb{R}_{\geq 0}$ is a transition **rate matrix** such that
  - $\mathbf{R}(s, \alpha, s') = \lambda$ is the rate of a negative exponential distribution

$$f_X(t) = \begin{cases} \lambda \cdot e^{-\lambda \cdot t} & \text{if } t \geq 0 \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad E[X] = \frac{1}{\lambda}$$

such that $Act(s) = \{\alpha \in Act \mid \exists s' \in \mathcal{S}.\ \mathbf{R}(s, \alpha, s') > 0\} \neq \emptyset$ for all $s \in \mathcal{S}$.

  - $E(s, \alpha) = \sum_{s' \in \mathcal{S}} \mathbf{R}(s, \alpha, s')$ is the **exit rate** of s under α.

Example

1. Nondeterministically choose $\beta \in Act(s_0)$.
2. Race between δ-transitions in s_2:
   - **Mean delay**: $\frac{1}{E(s_2,\delta)} = 4$.
   - **Probability** to move to s_4: $\frac{\mathbf{R}(s_2,\delta,s_4)}{E(s_4,\delta)} = \frac{8}{9}$.

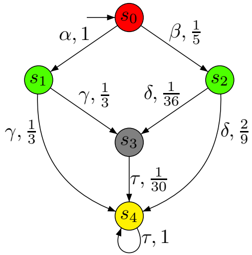

Trajectories in CTMDPs

1. **Finite paths** of length $n \in \mathbb{N}$ are denoted $\pi = s_0 \xrightarrow{\alpha_0,t_0} \cdots \xrightarrow{\alpha_{n-1},t_{n-1}} s_n$.
   - $\pi{\downarrow} = s_n$ is the last state of π.
   - $Paths^n$ is the set of paths of length n and
2. $Paths^\omega$ is the set of **infinite paths**.

Trajectories in CTMDPs

1. **Finite paths** of length $n \in \mathbb{N}$ are denoted $\pi = s_0 \xrightarrow{\alpha_0, t_0} \cdots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n$.
   - $\pi{\downarrow} = s_n$ is the last state of π.
   - $Paths^n$ is the set of paths of length n and
2. $Paths^\omega$ is the set of **infinite paths**.

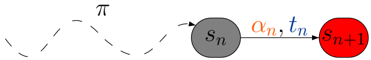

A **combined transition** $m = (\alpha_n, t_n, s_{n+1})$:

- α_n is the action in state $\pi{\downarrow}$ (chosen externally),
- t_n is the transition's firing time and
- s_{n+1} the transition's successor state.

$$\Omega := Act \times \mathbb{R}_{\geq 0} \times \mathcal{S} \text{ is the set of all combined transitions.}$$

Constructing events in CTMDPs

Probability measures are defined on σ-fields

1. $\mathfrak{F}$ of sets of combined transitions:

$$\Omega := Act \times \mathbb{R}_{\geq 0} \times \mathcal{S}$$

$$\mathfrak{F} := \sigma\big(\mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{\mathcal{S}}\big)$$

$\mathfrak{B}(\mathbb{R}_{\geq 0})$: Borel σ-field for $\mathbb{R}_{\geq 0}$

2. $\mathfrak{F}_{Paths^n}$ of sets of paths of length n:

$$\mathfrak{F}_{Paths^n} := \sigma\big(\{S_0 \times M_1 \times \cdots \times M_n \mid S_0 \in \mathfrak{F}_{\mathcal{S}}, M_i \in \mathfrak{F}\}\big)$$

3. $\mathfrak{F}_{Paths^\omega}$ of sets of infinite paths:

Cylinder set construction:

- Any $C^n \in \mathfrak{F}_{Paths^n}$ defines a cylinder base (of finite length)
- $C_n := \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$ is a cylinder (extension to infinity).

The σ-field $\mathfrak{F}_{Paths^\omega}$ is then

$$\mathfrak{F}_{Paths^\omega} := \sigma\Big(\bigcup_{n=0}^{\infty} \{C_n \mid C^n \in \mathfrak{F}_{Paths^n}\}\Big)$$

Constructing events in CTMDPs

Probability measures are defined on σ-fields

1. $\mathfrak{F}$ of sets of combined transitions:

$$\Omega := Act \times \mathbb{R}_{\geq 0} \times \mathcal{S}$$

$$\mathfrak{F} := \sigma\big(\mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{\mathcal{S}}\big)$$

$\mathfrak{B}(\mathbb{R}_{\geq 0})$: Borel σ-field for $\mathbb{R}_{\geq 0}$

2. $\mathfrak{F}_{Paths^n}$ of sets of paths of length n:

$$\mathfrak{F}_{Paths^n} := \sigma\left(\{S_0 \times M_1 \times \cdots \times M_n \mid S_0 \in \mathfrak{F}_{\mathcal{S}}, M_i \in \mathfrak{F}\}\right)$$

3. $\mathfrak{F}_{Paths^\omega}$ of sets of infinite paths:

Cylinder set construction:

- Any $C^n \in \mathfrak{F}_{Paths^n}$ defines a cylinder base (of finite length)
- $C_n := \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$ is a cylinder (extension to infinity).

The σ-field $\mathfrak{F}_{Paths^\omega}$ is then

$$\mathfrak{F}_{Paths^\omega} := \sigma\Big(\bigcup_{n=0}^{\infty} \{C_n \mid C^n \in \mathfrak{F}_{Paths^n}\}\Big)$$

Constructing events in CTMDPs

Probability measures are defined on σ-fields

1. $\mathfrak{F}$ of sets of combined transitions:

$$\Omega := Act \times \mathbb{R}_{\geq 0} \times \mathcal{S}$$

$$\mathfrak{F} := \sigma\big(\mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{\mathcal{S}}\big)$$

$\mathfrak{B}(\mathbb{R}_{\geq 0})$: Borel σ-field for $\mathbb{R}_{\geq 0}$

2. $\mathfrak{F}_{Paths^n}$ of sets of paths of length n:

$$\mathfrak{F}_{Paths^n} := \sigma\left(\{S_0 \times M_1 \times \cdots \times M_n \mid S_0 \in \mathfrak{F}_{\mathcal{S}}, M_i \in \mathfrak{F}\}\right)$$

3. $\mathfrak{F}_{Paths^\omega}$ of sets of infinite paths:

Cylinder set construction:

- Any $C^n \in \mathfrak{F}_{Paths^n}$ defines a cylinder base (of finite length)
- $C_n := \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$ is a **cylinder** (extension to infinity).

The σ-field $\mathfrak{F}_{Paths^\omega}$ is then

$$\mathfrak{F}_{Paths^\omega} := \sigma\Big(\bigcup_{n=0}^{\infty} \{C_n \mid C^n \in \mathfrak{F}_{Paths^n}\}\Big)$$

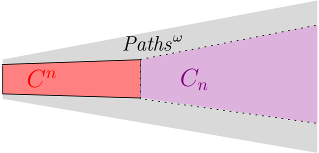

The probability of events

Resolving nondeterminism: Assume state s_n is hit after trajectory

$$\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} s_2 \xrightarrow{\alpha_2, t_2} \cdots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n.$$

- Nondeterminism occurs in s_n if $|Act(s_n)| > 1$.
- A scheduler resolves it and uniquely induces a stochastic process.

The probability of events

Resolving nondeterminism: Assume state s_n is hit after trajectory

$$\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} s_2 \xrightarrow{\alpha_2, t_2} \cdots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n.$$

- Nondeterminism occurs in s_n if $|Act(s_n)| > 1$.
- A scheduler resolves it and uniquely induces a stochastic process.

A hierarchy of scheduler classes:

1. Generic measurable scheduler (GM): $D : Paths^{\star} \rightarrow Distr(Act)$
2. Total time positional scheduler (TTP): $D : \mathcal{S} \times \mathbb{R}_{\geq 0} \rightarrow Distr(Act)$
3. Time abstract hop counting scheduler (TAHOP): $D : \mathcal{S} \times \mathbb{N} \rightarrow Distr(Act)$
4. Time abstract positional scheduler (TAP): $D : \mathcal{S} \rightarrow Distr(Act)$

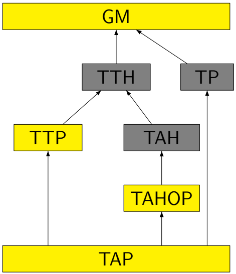

The probability of a single step $M \subseteq \mathfrak{F}$

1. Enter state s_n along trajectory
   $\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \cdots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n$.
2. Continue in s_n with a transition
   $(\alpha_n, t_n, s_{n+1}) \in M$
3. Measure probability of sets $M \subseteq \mathfrak{F}$!
   Example: $M = \{\alpha_n\} \times [0,1] \times \{s_{n+1}\}$.

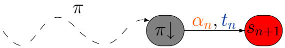

Probability measure $\mu_D(\pi, \cdot) : \mathfrak{F} \to [0,1]$ on sets of combined transitions:

- Choose an action, wait and jump to successor state.

$$\mu_D(\pi, M) := \int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow, \alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \; \mathbf{P}(\pi\downarrow, \alpha, ds').$$

- Note: $\eta_{E(\pi\downarrow, \alpha)}$ depends on scheduler D!
  Therefore: Scheduler cannot incorporate the sojourn time in state $\pi\downarrow$.

A generic probability measure on sets of paths

1. Initial distribution ν: Probability to start in state s.

2. $Pr^n_{\nu,D}$ on sets of finite paths:
Let $\nu \in Distr(S)$ and $D \in TTP$. Define inductively:

$$Pr^0_{\nu,D}(\Pi) := \sum_{s \in \Pi} \nu(s) \quad \text{and for } n > 0$$

$$Pr^n_{\nu,D}(\Pi) := \int_{Paths^{n-1}} Pr^{n-1}_{\nu,D}(d\pi) \int_{\Omega} \mathbf{1}_{\Pi}(\pi \circ m) \; \mu_D(\pi, dm).$$

3. $Pr^\omega_{\nu,D}$ on sets of infinite paths:
   - A cylinder base is a measurable set $C^n \in \mathfrak{F}_{Paths^n}$
   - C^n defines cylinder $C_n = \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$
   - The probability of cylinder C_n is that of its base C^n:

$$Pr^\omega_{\nu,D}(C_n) = Pr^n_{\nu,D}(C^n).$$

   This extends to $\mathfrak{F}_{Paths^\omega}$ by Ionescu-Tulcea.

A generic probability measure on sets of paths

1. Initial distribution ν: Probability to start in state s.
2. $Pr^n_{\nu,D}$ on sets of finite paths:
   Let $\nu \in Distr(\mathcal{S})$ and $D \in TTP$. Define inductively:

$$Pr^0_{\nu,D}(\Pi) := \sum_{s \in \Pi} \nu(s) \quad \text{and for } n > 0$$

$$Pr^n_{\nu,D}(\Pi) := \int_{Paths^{n-1}} Pr^{n-1}_{\nu,D}(d\pi) \int_{\Omega} \mathbf{I}_{\Pi}(\pi \circ m) \; \mu_D(\pi, dm) \, .$$

3. $Pr^\omega_{\nu,D}$ on sets of infinite paths:
   - A cylinder base is a measurable set $C^n \in \mathfrak{F}_{Paths^n}$
   - C^n defines cylinder $C_n = \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$
   - The probability of cylinder C_n is that of its base C^n:

$$Pr^\omega_{\nu,D}(C_n) = Pr^n_{\nu,D}(C^n).$$

   This extends to $\mathfrak{F}_{Paths^\omega}$ by Ionescu-Tulcea.

A generic probability measure on sets of paths

1. Initial distribution ν: Probability to start in state s.
2. $Pr^n_{\nu,D}$ on sets of finite paths:
   Let $\nu \in Distr(\mathcal{S})$ and $D \in TTP$. Define inductively:

$$Pr^0_{\nu,D}(\Pi) := \sum_{s \in \Pi} \nu(s) \quad \text{and for } n > 0$$

$$Pr^n_{\nu,D}(\Pi) := \int_{Paths^{n-1}} Pr^{n-1}_{\nu,D}(d\pi) \int_{\Omega} \mathbf{I}_{\Pi}(\pi \circ m) \; \mu_D(\pi, dm).$$

3. $Pr^\omega_{\nu,D}$ on sets of infinite paths:
   - A **cylinder base** is a measurable set $C^n \in \mathfrak{F}_{Paths^n}$
   - C^n defines **cylinder** $C_n = \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$
   - The probability of cylinder C_n is that of its base C^n:

$$Pr^\omega_{\nu,D}(C_n) = Pr^n_{\nu,D}(C^n).$$

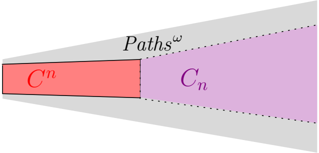

This extends to $\mathfrak{F}_{Paths^\omega}$ by **Ionescu-Tulcea**.

Delaying the resolution of nondeterminism

- The semantics of a single step so far:

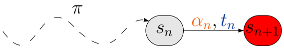

1. Scheduler decides upon entering s_n.
2. Sojourn time in s_n depends on choice!

$$\int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow,\alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds')$$

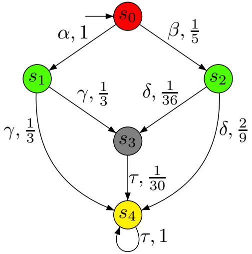

Delaying the resolution of nondeterminism

- The semantics of a single step so far:

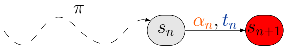

  1. Scheduler decides upon entering s_n.
  2. Sojourn time in s_n depends on choice!

$$\int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow,\alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \ \mathbf{P}(\pi\downarrow, \alpha, ds')$$

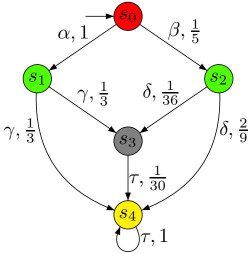

- Idea to delay resolution of nondeterminism:
  Schedule only when the current state is left!
  Therefore: Dissolve dependency between
  - sojourn time in state s_n and
  - scheduler's choice when entering s_n.

$$\int_{\mathbb{R}_{\geq 0}} \eta_{\lambda(s_n)}(dt) \int_{Act} D(\pi, t, d\alpha) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \ \mathbf{P}(\pi\downarrow, \alpha, ds')$$

Delaying the resolution of nondeterminism

- The semantics of a single step so far:

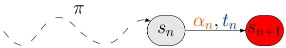

  1. Scheduler decides upon entering s_n.
  2. Sojourn time in s_n depends on choice!

$$\int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow,\alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds')$$

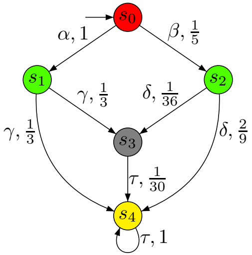

- Idea to delay resolution of nondeterminism:
  Schedule only when the current state is left!
  Therefore: Dissolve dependency between
  - sojourn time in state s_n and
  - scheduler's choice when entering s_n.

$$\int_{\mathbb{R}_{\geq 0}} \eta_{\lambda(s_n)}(dt) \int_{Act} D(\pi, t, d\alpha) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds')$$

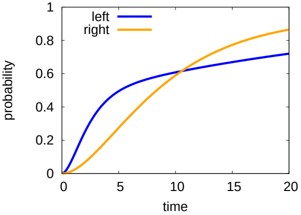

Local uniformity enables delayed scheduling

non-uniform CTMDP

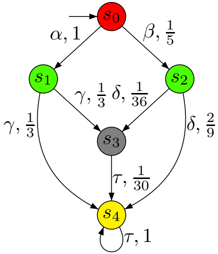

A CTMDP $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, \nu)$ is **locally uniform** iff there exists $\lambda : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ s.t.

$$\forall s \in \mathcal{S}. \, \forall \alpha \in Act(s). \quad \lambda(s) = E(s, \alpha).$$

Local uniformity enables delayed scheduling

A CTMDP $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, \nu)$ is **locally uniform** iff there exists $\lambda : \mathcal{S} \to \mathbb{R}_{\geq 0}$ s.t.

$$\forall s \in \mathcal{S}.\ \forall \alpha \in Act(s). \quad \lambda(s) = E(s, \alpha).$$

non-uniform CTMDP

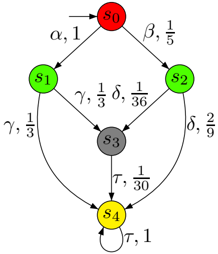

local uniformization

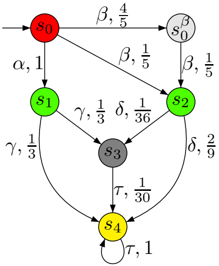

Local uniformity enables delayed scheduling

non-uniform CTMDP

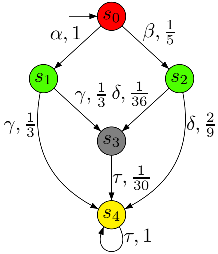

A CTMDP $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, \nu)$ is **locally uniform** iff there exists $\lambda : \mathcal{S} \to \mathbb{R}_{\geq 0}$ s.t.

$$\forall s \in \mathcal{S}. \forall \alpha \in Act(s). \quad \lambda(s) = E(s, \alpha).$$

Local uniformization yields $unif(\mathcal{C}) = (\overline{\mathcal{S}}, Act, \overline{\mathbf{R}}, \nu)$:

- $\overline{\mathcal{S}} = \mathcal{S} \uplus \{s^\alpha \mid s \in \mathcal{S}, \alpha \in Act$ with $E(s, \alpha) < \lambda(s)\}$
- $\overline{\mathbf{R}}(s, \alpha, s') = \begin{cases} \mathbf{R}(s, \alpha, s') & \text{if } s, s' \in \mathcal{S} \\ \lambda(s) - E(s, \alpha) & \text{if } s \in \mathcal{S} \text{ and } s' = s^\alpha \\ \mathbf{R}(t, \alpha, s') & \text{if } s = t^\alpha \text{ and } s' \in \mathcal{S} \\ 0 & \text{otherwise.} \end{cases}$

local uniformization

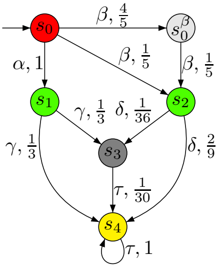

A hint towards correctness of local uniformization

non-uniform CTMDP

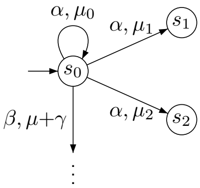

$E(s, \alpha) = \mu$ and $E(s, \beta) = \mu + \gamma$

A hint towards correctness of local uniformization

non-uniform CTMDP

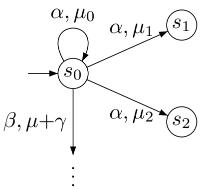

$E(s,\alpha) = \mu$ and $E(s,\beta) = \mu + \gamma$

locally uniform CTMDP

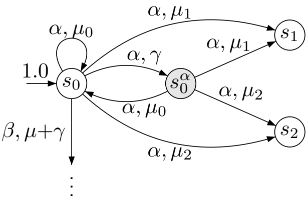

$E(s,\alpha) = E(s,\beta) = \mu + \gamma$

A hint towards correctness of local uniformization

non-uniform CTMDP

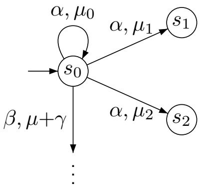

$E(s, \alpha) = \mu$ and $E(s, \beta) = \mu + \gamma$

locally uniform CTMDP

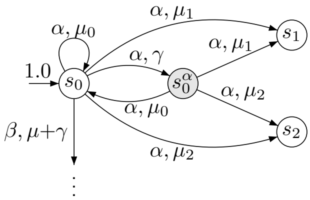

$E(s, \alpha) = E(s, \beta) = \mu + \gamma$

Correctness: If α is chosen in s, reachability of state u_i within $[0, t]$ is preserved:

$$\frac{\mu_i}{\mu} \int_0^t \eta_\mu(dt) = \frac{\mu_i}{\mu + \gamma} \int_0^t \eta_{\mu+\gamma}(dt_1) + \frac{\mu}{\mu + \gamma} \int_0^t \eta_{\mu+\gamma}(dt_1) \, \frac{\mu_i}{\mu} \int_0^{t-t_1} \eta_\mu(dt_2)$$

where $\eta_x = x \cdot e^{-x \cdot t}$ and $\mu = \sum \mu_i$.

A hint towards correctness of local uniformization

non-uniform CTMDP

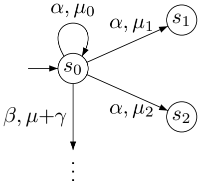

$E(s, \alpha) = \mu$ and $E(s, \beta) = \mu + \gamma$

locally uniform CTMDP

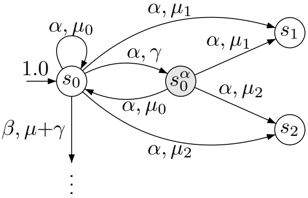

$E(s, \alpha) = E(s, \beta) = \mu + \gamma$

Correctness: If α is chosen in s, reachability of state u_i within $[0, t]$ is preserved:

$$\frac{\mu_i}{\mu} \int_0^t \eta_\mu(dt) = \frac{\mu_i}{\mu + \gamma} \int_0^t \eta_{\mu+\gamma}(dt_1) + \frac{\mu}{\mu + \gamma} \int_0^t \eta_{\mu+\gamma}(dt_1)\, \frac{\mu_i}{\mu} \int_0^{t-t_1} \eta_\mu(dt_2)$$

where $\eta_x = x \cdot e^{-x \cdot t}$ and $\mu = \sum \mu_i$.

But: No nondeterminism considered yet!

A correspondence between paths in $\mathcal{C}$ and $unif(\mathcal{C})$

non-uniform CTMDP

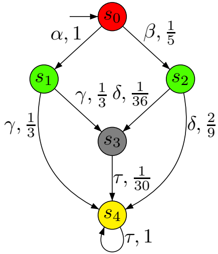

local uniformization

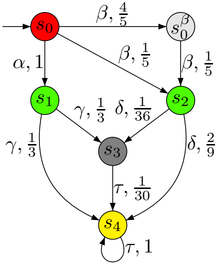

The function $merge : Paths(\overline{\mathcal{C}}) \to Paths(\mathcal{C})$ collapses copy-states s^α:

$$\overline{\pi} = s_0 \xrightarrow{\beta, t_0} s_0^\beta \xrightarrow{\beta, t_0'} s_2 \xrightarrow{\delta, t_1} s_4$$

$$merge(\overline{\pi}) = s_0 \xrightarrow{\beta, t_0 + t_0'} s_2 \xrightarrow{\delta, t_1} s_4.$$

The function $extend : Paths(\mathcal{C}) \to \mathfrak{F}_{Paths(\overline{\mathcal{C}})}$ is the inverse of $merge$.

Resolving nondeterminism in $unif(\mathcal{C})$

Any CTMDP $\mathcal{C}$ with GM scheduler D induces the measure $Pr^{\omega}_{\nu,D}$.

Question: **How to mimic D's behaviour on $unif(\mathcal{C})$ to obtain the same probability?**

Resolving nondeterminism in $unif(\mathcal{C})$

Any CTMDP $\mathcal{C}$ with GM scheduler D induces the measure $Pr^{\omega}_{\nu,D}$.

Question: **How to mimic D's behaviour on $unif(\mathcal{C})$ to obtain the same probability?**

Definition (stutter scheduler)

Let D be a GM scheduler on $\mathcal{C}$.
Define the stutter scheduler $\overline{D}$ on $unif(\mathcal{C})$:

$$\overline{D}(\overline{\pi}) := \begin{cases} D(\pi) & \text{if } \overline{\pi}\!\downarrow \in \mathcal{S} \wedge merge(\overline{\pi}) = \pi, \\ \{\alpha \mapsto 1\} & \text{if } \pi\!\downarrow = s^{\alpha}. \end{cases}$$

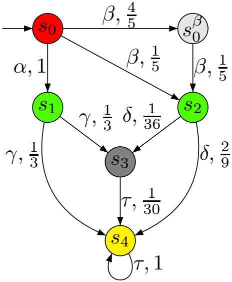

Resolving nondeterminism in $unif(\mathcal{C})$

Any CTMDP $\mathcal{C}$ with GM scheduler D induces the measure $Pr^{\omega}_{\nu,D}$.

Question: **How to mimic D's behaviour on $unif(\mathcal{C})$ to obtain the same probability?**

Definition (stutter scheduler)

Let D be a GM scheduler on $\mathcal{C}$.
Define the stutter scheduler $\overline{D}$ on $unif(\mathcal{C})$:

$$\overline{D}(\overline{\pi}) := \begin{cases} D(\pi) & \text{if } \overline{\pi}{\downarrow} \in \mathcal{S} \wedge merge(\overline{\pi}) = \pi, \\ \{\alpha \mapsto 1\} & \text{if } \pi{\downarrow} = s^{\alpha}. \end{cases}$$

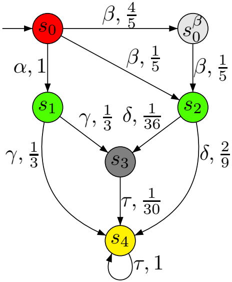

Note: No choice in copy-state s_0^{β}

Soundness: From $\mathcal{C}$ to $unif(\mathcal{C})$

The construction of $\overline{D}$ preserves all measures.

Proof sketch:

1. Uniformization is measure-preserving for measurable rectangles C^n:

$$Pr^n_{\nu,D}(C^n) = \overline{Pr}^{n}_{\nu,\overline{D}}(extend(C_n))$$

2. This extends to the field $\mathfrak{G}_{Paths^n} = \left(\mathfrak{F}_S \times \mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0})\right)^n \times \mathfrak{F}_S$.
3. Further we prove that

$$\mathfrak{C} = \left\{\Pi \in \mathfrak{F}_{Paths^n(\mathcal{C})} \mid Pr^n_{\nu,D}(\Pi) = \overline{Pr}^{n}_{\nu,\overline{D}}(extend(\Pi))\right\}$$

is a monotone class.

Soundness: From $\mathcal{C}$ to $unif(\mathcal{C})$

The construction of $\overline{D}$ preserves all measures.

Proof sketch:

1. Uniformization is measure-preserving for measurable rectangles C^n:

$$Pr^n_{\nu,D}(C^n) = \overline{Pr}^{\omega}_{\overline{\nu},\overline{D}}\big(extend(C_n)\big)$$

2. This extends to the field $\mathfrak{G}_{Paths^n} = \big(\mathfrak{F}_{\mathcal{S}} \times \mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0})\big)^n \times \mathfrak{F}_{\mathcal{S}}$.
3. Further we prove that

$$\mathfrak{C} = \Big\{\Pi \in \mathfrak{F}_{Paths^n(\mathcal{C})} \mid Pr^n_{\nu,D}(\Pi) = \overline{Pr}^{\omega}_{\overline{\nu},\overline{D}}\big(extend\big(\Pi\big)\big)\Big\}$$

is a monotone class.

Soundness: From $\mathcal{C}$ to $unif(\mathcal{C})$

The construction of $\overline{D}$ preserves all measures.

Proof sketch:

1. Uniformization is measure-preserving for measurable rectangles C^n:

$$Pr^n_{\nu,D}(C^n) = \overline{Pr}^{\omega}_{\overline{\nu},\overline{D}}\big(extend(C_n)\big)$$

2. This extends to the field $\mathfrak{G}_{Paths^n} = \big(\mathfrak{F}_{\mathcal{S}} \times \mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0})\big)^n \times \mathfrak{F}_{\mathcal{S}}$.
3. Further we prove that

$$\mathfrak{C} = \Big\{\Pi \in \mathfrak{F}_{Paths^n(\mathcal{C})} \mid Pr^n_{\nu,D}(\Pi) = \overline{Pr}^{\omega}_{\overline{\nu},\overline{D}}\big(extend\big(\Pi\big)\big)\Big\}$$

is a monotone class.

The claim follows by applying the **Monotone Class Theorem**.

Completeness: From $unif(\mathcal{C})$ to $\mathcal{C}$.

Main results:

1. For scheduler classes $\mathfrak{G} \in \{TTP, TAP\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) = \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

2. For the classes $\mathfrak{G} \in \{TAHOP, TAH, TP\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) \neq \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

3. Our main concern: Timed reachability analysis:
   - Previous results hold for arbitrary measures.
   - Reachability of states in G in time t:

$$\sup_{D \in TTP(\mathcal{C})} Pr^{\omega}_{\nu,D}\left(\Diamond^{[0,t]} G\right) = \sup_{D \in GM(\mathcal{C})} Pr^{\omega}_{\nu,D}\left(\Diamond^{[0,t]} G\right)$$

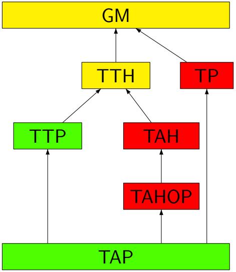

Completeness: From $unif(\mathcal{C})$ to $\mathcal{C}$.

Main results:

1. For scheduler classes $\mathfrak{G} \in \{TTP, TAP, TTH, GM\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) = \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

Conjecture: GM and TTH are also complete.

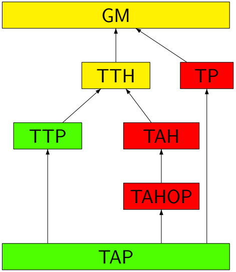

Completeness: From $unif(\mathcal{C})$ to $\mathcal{C}$.

Main results:

1. For scheduler classes $\mathfrak{G} \in \{TTP, TAP, TTH, GM\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) = \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

Conjecture: GM and TTH are also complete.

2. For the classes $\mathfrak{G} \in \{TAHOP, TAH, TP\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) \neq \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

3. Our main concern: Timed reachability analysis:
   - Previous results hold for arbitrary measures.
   - Reachability of states in G in time t:

$$\sup_{D \in TTP(\mathcal{C})} Pr^{\omega}_{\nu,D}\left(\Diamond^{[0,t]} G\right) = \sup_{D \in GM(\mathcal{C})} Pr^{\omega}_{\nu,D}\left(\Diamond^{[0,t]} G\right).$$

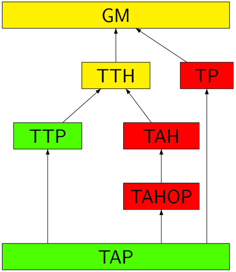

Completeness: From $unif(\mathcal{C})$ to $\mathcal{C}$.

Main results:

1. For scheduler classes $\mathfrak{G} \in \{TTP, TAP, TTH, GM\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) = \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

Conjecture: GM and TTH are also complete.

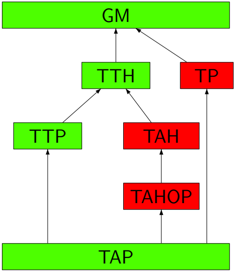

2. For the classes $\mathfrak{G} \in \{TAHOP, TAH, TP\}$:

$$\sup_{D \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D}(\Pi) \neq \sup_{D' \in \mathfrak{G}(\mathcal{C})} Pr^{\omega}_{\nu,D'}(extend(\Pi))$$

3. Our main concern: **Timed reachability analysis**:
   - Previous results hold for arbitrary measures.
   - Reachability of states in G in time t:

$$\sup_{D \in TTP(\mathcal{C})} Pr^{\omega}_{\nu,D}\left(\Diamond^{[0,t]} G\right) = \sup_{D \in GM(\mathcal{C})} Pr^{\omega}_{\nu,D}\left(\Diamond^{[0,t]} G\right).$$

The benefit of delaying nondeterminism

- Instead of **early scheduling**:

$$\mu_D(\pi, M) = \int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow,\alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds'),$$

- local uniformity allows **late scheduling**:

$$\mu_D(\pi, M) = \int_{\mathbb{R}_{\geq 0}} \eta_{\lambda(s_n)}(dt) \int_{Act} D(\pi, t, d\alpha) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds').$$

- What's the benefit?

The benefit of delaying nondeterminism

- Instead of **early scheduling**:

$$\mu_D(\pi, M) = \int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow,\alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \ \ \mathbf{P}(\pi\downarrow, \alpha, ds'),$$

- local uniformity allows **late scheduling**:

$$\mu_D(\pi, M) = \int_{\mathbb{R}_{\geq 0}} \eta_{\lambda(s_n)}(dt) \int_{Act} D(\pi, t, d\alpha) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \ \ \mathbf{P}(\pi\downarrow, \alpha, ds').$$

- What's the benefit?

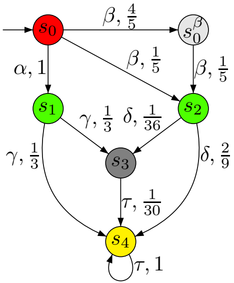

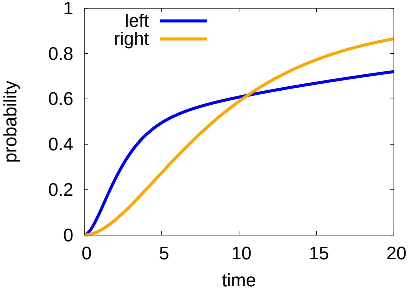

The benefit of delaying nondeterminism

- Instead of **early scheduling**:

$$\mu_D(\pi, M) = \int_{Act} D(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi\downarrow,\alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds'),$$

- local uniformity allows **late scheduling**:

$$\mu_D(\pi, M) = \int_{\mathbb{R}_{\geq 0}} \eta_{\lambda(s_n)}(dt) \int_{Act} D(\pi, t, d\alpha) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s') \;\; \mathbf{P}(\pi\downarrow, \alpha, ds').$$

- What's the benefit?

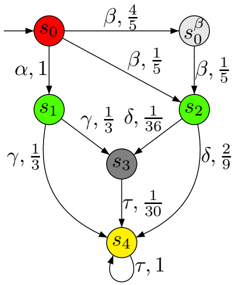

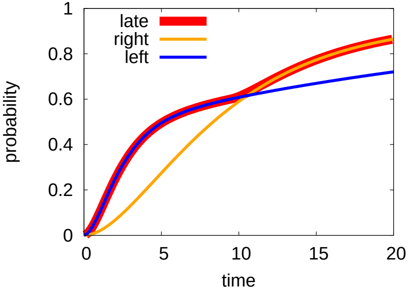

What is achieved:

We consider **locally uniform CTMDPs** and **late schedulers**:

1. They allow to delay the resolution of nondeterminism.
2. Late schedulers are strictly better than any early scheduler.

What is achieved:

We consider **locally uniform CTMDPs** and **late schedulers**:

1. They allow to delay the resolution of nondeterminism.
2. Late schedulers are strictly better than any early scheduler.

We investigate a **transformation** which achieves local uniformity.

1. Local uniformization works for important scheduler classes.
2. The transform is viable to late scheduling.
3. **Future work: Timed reachability analysis.**
   Late scheduling has proved to be **algorithmically manageable!**

What is achieved:

We consider **locally uniform CTMDPs** and **late schedulers**:

1. They allow to delay the resolution of nondeterminism.
2. Late schedulers are strictly better than any early scheduler.

We investigate a **transformation** which achieves local uniformity.

1. Local uniformization works for important scheduler classes.
2. The transform is viable to late scheduling.
3. **Future work: Timed reachability analysis.**
   Late scheduling has proved to be **algorithmically manageable!**

Thank you for your attention!