Deciding the value 1 problem for reachability in 1-clock Decision Stochastic Timed Automata

Nathalie Bertrand[1] Thomas Brihaye[2] Blaise Genest[3]

[1] Inria/IRISA, Rennes, France

[2] Université Mons, Mons, Belgium

[3] CNRS/IRISA, Rennes, France

QEST'14 - 9/9 - Firenze

Non-deterministic and probabilistic timed systems

Two approaches to combine probability, non-determinism and time:

- Probabilistic Timed Automata (*à la* PRISM) [KNSS-arts99]
  - Time-delays are chosen non-deterministically,
  - Edges are according to discrete probability distributions.
- Decision Stochastic Timed Automata (ext. of CTMDP) [BS-formats12]
  - Time-delays are chosen via continuous probability distributions,
  - Edges are chosen non-deterministically.

Known results on (decision) stochastic timed automata

Stochastic timed automata (STA)

- The almost-sure model-checking of LTL is **decidable**
  - on 1-clock STA. [BBBBG-lics08]
  - on reactive n-clock STA. [BBJM-qest12]
- **Open problem** decidability of the almost-sure reachability problem on general 2-clock STA.

Decision Stochastic Timed Automata (DSTA)

- **Existence** of an optimal scheduler for the **time-bounded reachability problem** on reactive DSTA. [BS-formats12]

Known results on (decision) stochastic timed automata

Stochastic timed automata (STA)

- The almost-sure model-checking of LTL is **decidable**
  - on 1-clock STA. [BBBBG-lics08]
  - on reactive n-clock STA. [BBJM-qest12]
- **Open problem** decidability of the almost-sure reachability problem on general 2-clock STA.

Decision Stochastic Timed Automata (DSTA)

- **Existence** of an optimal scheduler for the **time-bounded reachability problem** on reactive DSTA. [BS-formats12]

This talk: **reachability problem on 1-clock DSTA**.

Outline of the talk

One-clock timed automata

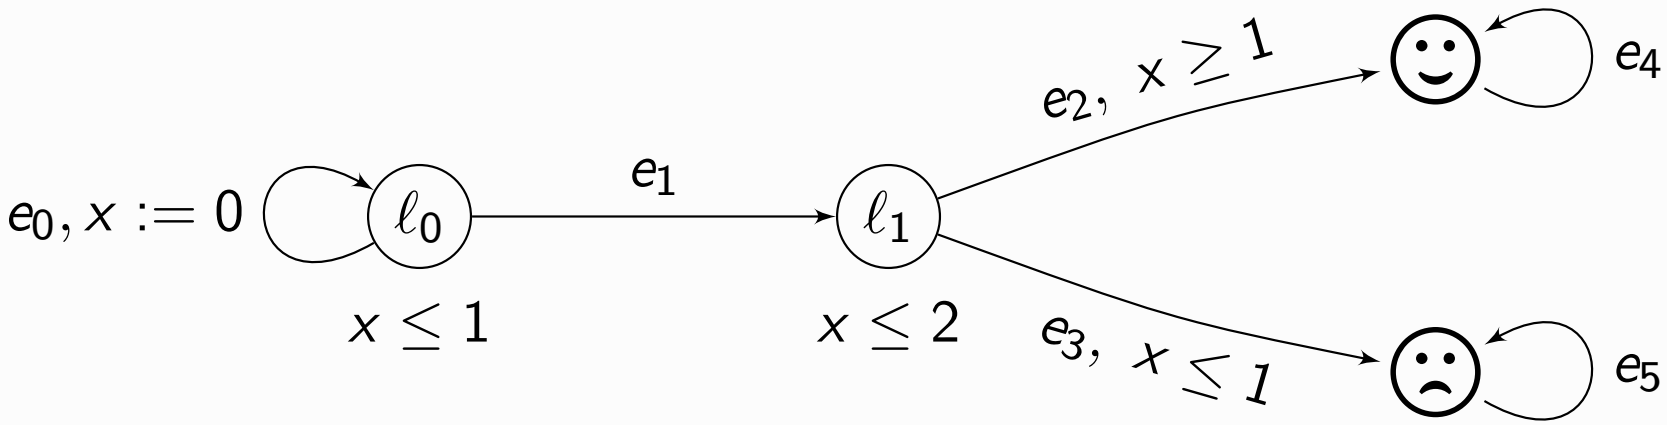

one-clock timed automaton: $\mathcal{A} = (L, \ell_0, E, \mathcal{I})$

Example of execution:

$$(\ell_0, 0) \xrightarrow{.7} (\ell_0, .7) \xrightarrow{e_0} (\ell_0, 0) \xrightarrow{.8} (\ell_0, .8) \xrightarrow{e_1} (\ell_1, .8) \xrightarrow{.3} (\ell_0, 1.1) \xrightarrow{e_2} \text{☺}$$

One-clock Decision Stochastic Timed Automaton

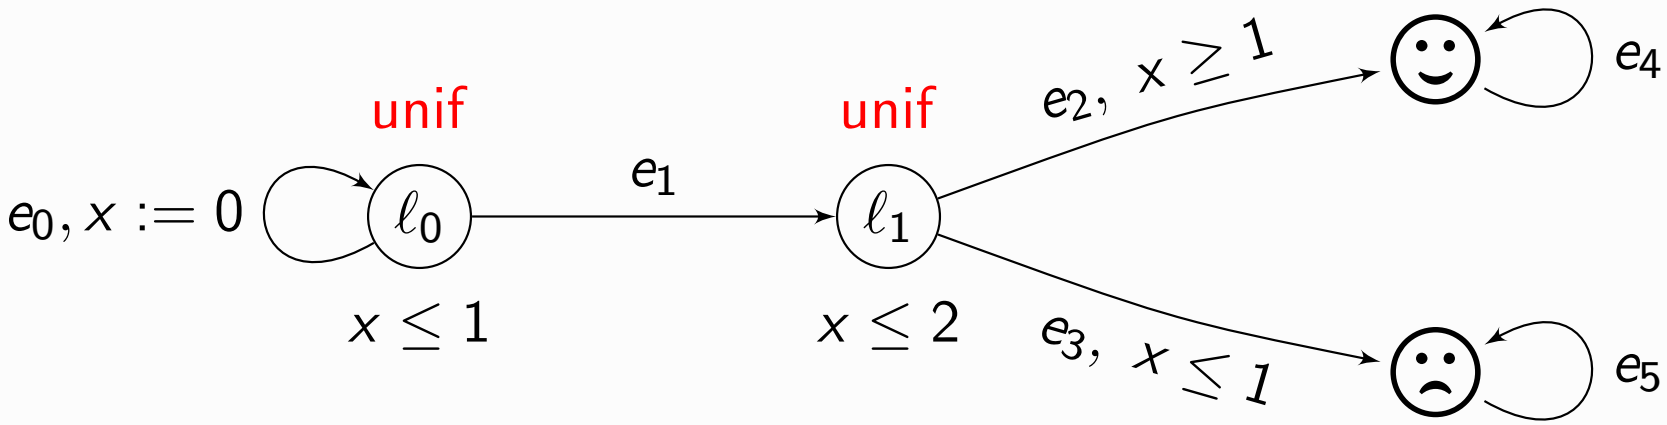

decision stochastic timed automaton: $(\mathcal{A}, \mu)$ where

- $\mathcal{A} = (L, \ell_0, E, \mathcal{I})$ is a one-clock timed automaton
- $\mu = (\mu_{\ell,\nu})$ is a family of *distributions*
  $\mu_{\ell,\nu}$: distribution over potential delays from state (ℓ, ν)

Semantics of DSTA

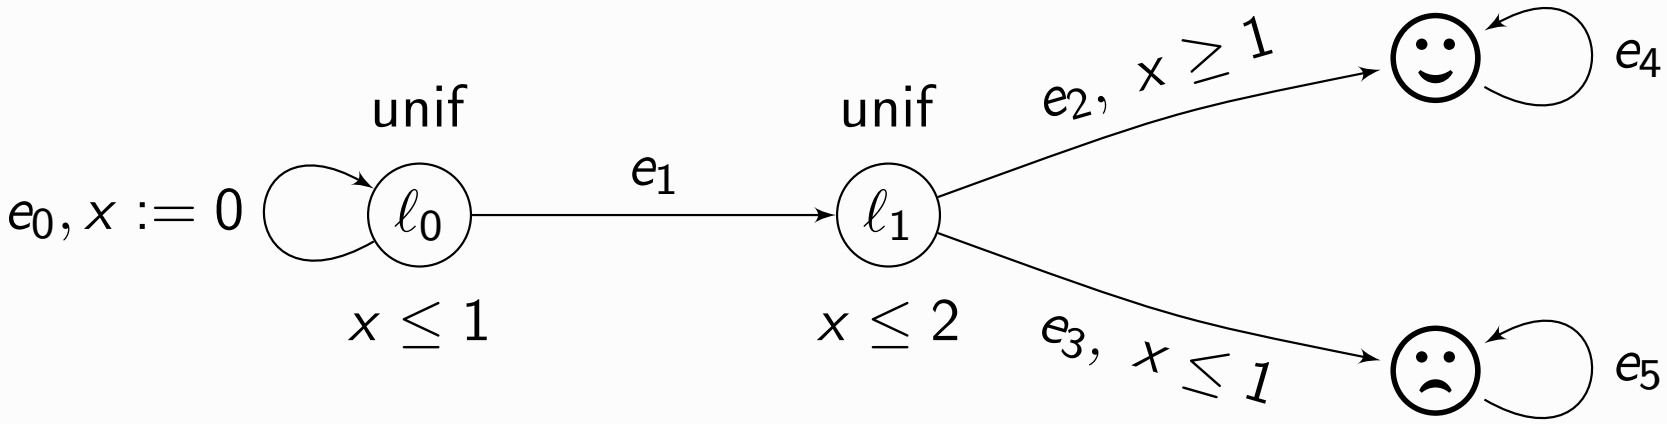

Infinite state MDP: from state s

- a delay τ is randomly chosen according to μ_s;
- the *player* chooses an edge e enabled in $s + \tau$.

$$\langle \ell_0, 0 \rangle \xrightarrow{.7} [\ell_0, .7] \xrightarrow{e_0} \langle \ell_0, 0 \rangle \xrightarrow{.8} [\ell_0, .8] \xrightarrow{e_1} \langle \ell_1, .8 \rangle \xrightarrow{.3} [\ell_0, 1.1] \xrightarrow{e_2} \text{☺}$$

strategy σ: in [ ]-states dictates which edge to choose

Optimal positional strategies

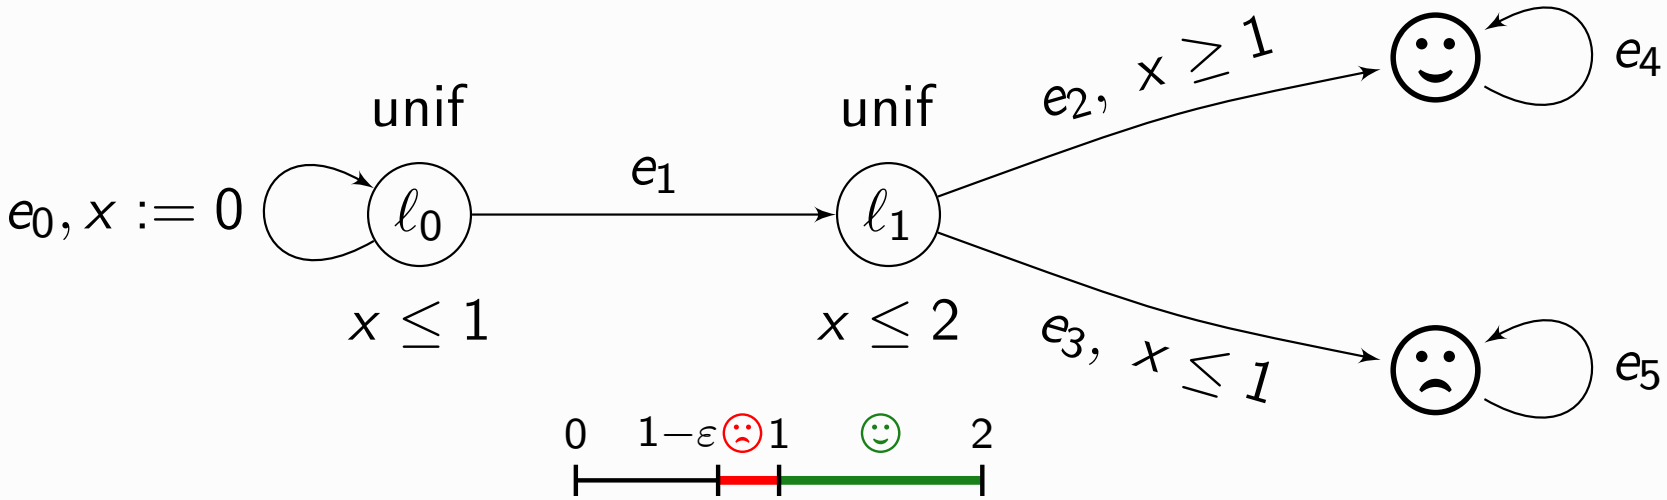

$$\sigma_\varepsilon(\ell_0, \nu) = \begin{cases} e_0 & \text{if } \nu < 1-\varepsilon \\ e_1 & \text{if } \nu \geq 1-\varepsilon \end{cases} \quad ; \quad \sigma_\varepsilon(\ell_1, \nu) = \begin{cases} e_2 & \text{if } \nu \geq 1 \\ e_3 & \text{if } \nu < 1 \end{cases}$$

$$\mathbb{P}^{(\ell_0,0)}_{\sigma_\varepsilon}\big((\mathcal{A},\mu) \models \Diamond \text{☺}\big) \geq \frac{1}{1+\varepsilon} \geq 1-\varepsilon.$$

Optimal positional strategies

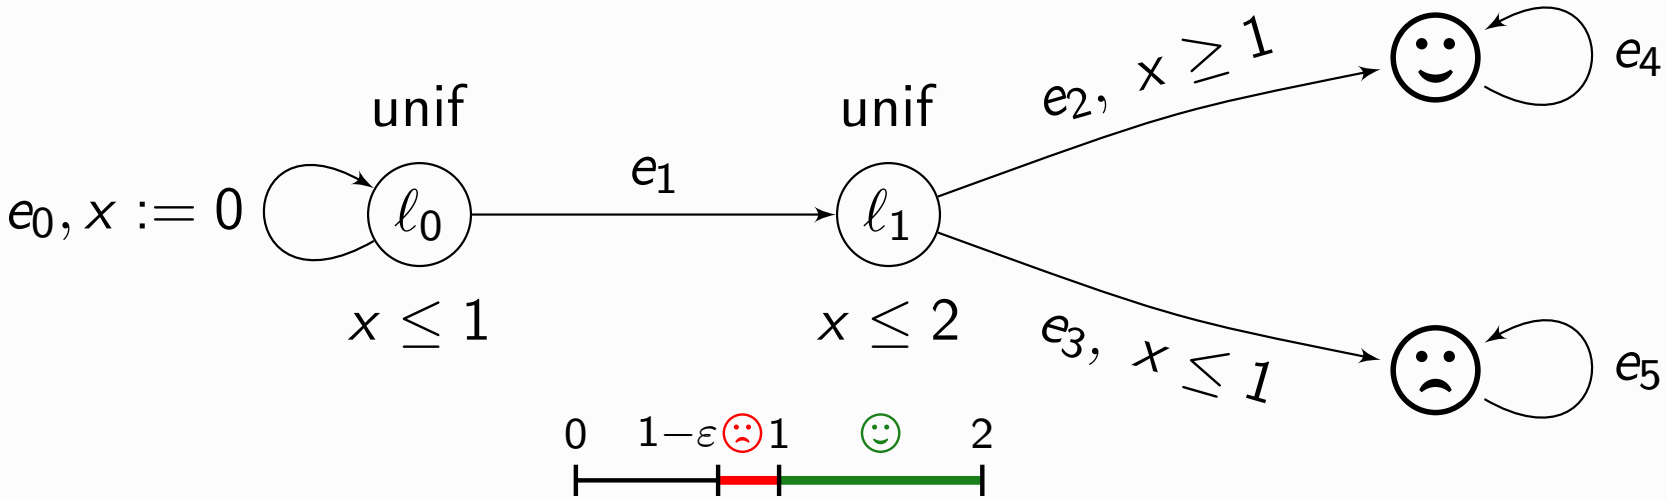

$$\sigma_\varepsilon(\ell_0, \nu) = \begin{cases} e_0 & \text{if } \nu < 1-\varepsilon \\ e_1 & \text{if } \nu \geq 1-\varepsilon \end{cases} \quad ; \quad \sigma_\varepsilon(\ell_1, \nu) = \begin{cases} e_2 & \text{if } \nu \geq 1 \\ e_3 & \text{if } \nu < 1 \end{cases}$$

$$\mathbb{P}^{(\ell_0,0)}_{\sigma_\varepsilon}\big((\mathcal{A},\mu) \models \Diamond \text{☺}\big) \geq \frac{1}{1+\varepsilon} \geq 1-\varepsilon.$$

ε-optimal strategies are not region-uniform

Almost-sure vs Limit-sure

DSTA $(\mathcal{A}, \mu)$, target set ☺ $\subseteq L$, and initial state $s \in S$

- ☺ is **almost-surely** reachable from s if

$$\exists \sigma \quad \mathbb{P}^{s}_{\sigma}\big((\mathcal{A}, \mu) \models \Diamond ☺\big) = 1.$$

- ☺ is **limit-surely** reachable from s if

$$\forall \varepsilon > 0 \; \exists \sigma \quad \mathbb{P}^{s}_{\sigma}\big((\mathcal{A}, \mu) \models \Diamond ☺\big) > 1 - \varepsilon.$$

Almost-sure $\neq$ Limit-sure

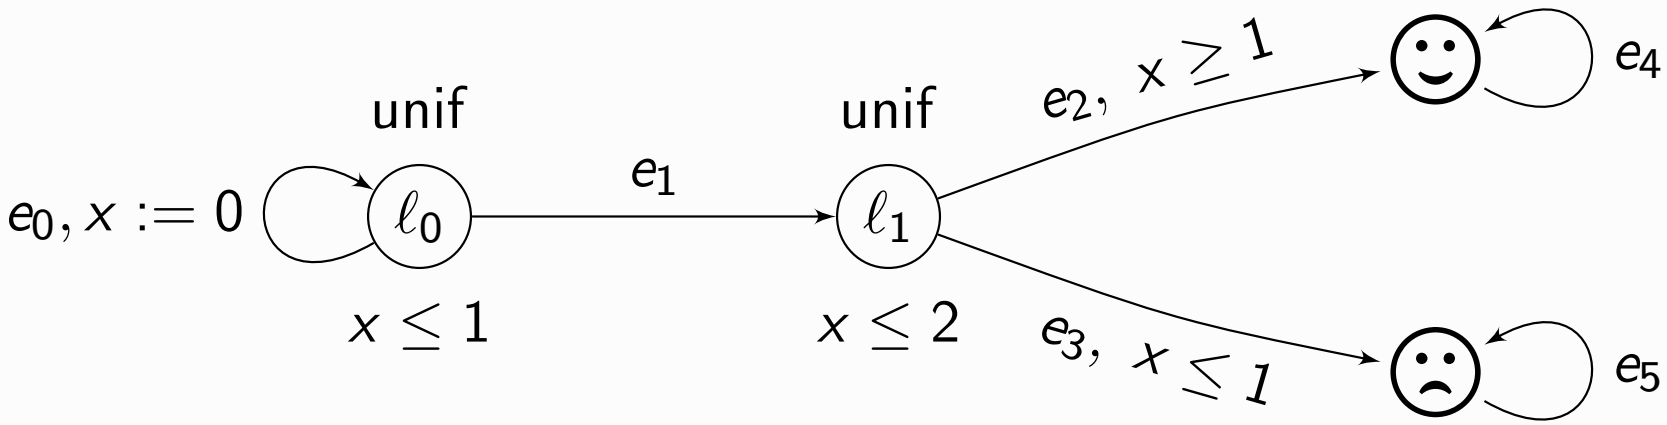

- ☺ is not almost-surely reachable from $(\ell_0, 0)$,
- ☺ is limit-surely reachable from $(\ell_0, 0)$.

Our contribution

Probability 1 problem

Input: A DSTA $(\mathcal{A}, \mu)$, a target set $\☺ \subseteq L$ and a state $s \in S$.
Question: Is $\☺$ almost-surely reachable from s?

Value 1 problem

Input: A DSTA $(\mathcal{A}, \mu)$, a target set $\☺ \subseteq L$ and a state $s \in S$.
Question: Is $\☺$ limit-surely reachable from s?

Our contribution

Probability 1 problem

Input: A DSTA $(\mathcal{A}, \mu)$, a target set $\☺ \subseteq L$ and a state $s \in S$.
Question: Is $\☺$ almost-surely reachable from s?

Value 1 problem

Input: A DSTA $(\mathcal{A}, \mu)$, a target set $\☺ \subseteq L$ and a state $s \in S$.
Question: Is $\☺$ limit-surely reachable from s?

Main Result

The probability 1 and value 1 problems are decidable in polynomial time for one-clock decision stochastic timed automata.
For value 1, ε-optimal strategies are not region-uniform.

Solving the value 1 problem - Key idea

From a DSTA $(\mathcal{A}, \mu)$, we build a **finite MDP** $\mathcal{A}_{\mathsf{cp}}$ such that

☺ is **limit-surely** reachable from s_0 in $(\mathcal{A}, \mu)$

if and only if

☺ is **almost-surely** reachable from s_0 in $\mathcal{A}_{\mathsf{cp}}$.

Limit corner-point region MDP

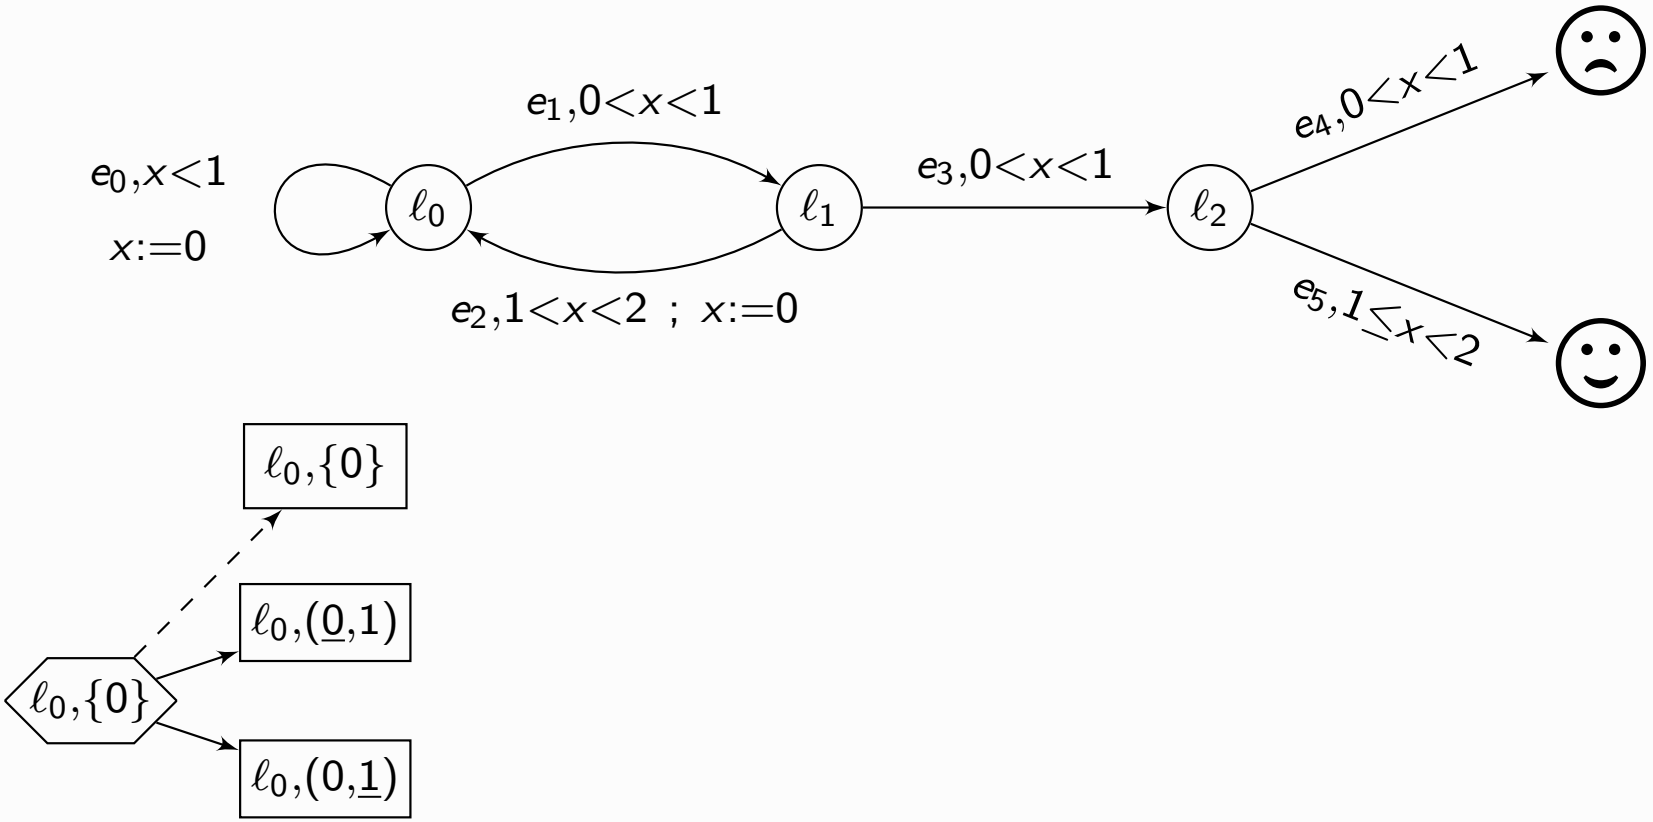

Limit corner-point region MDP

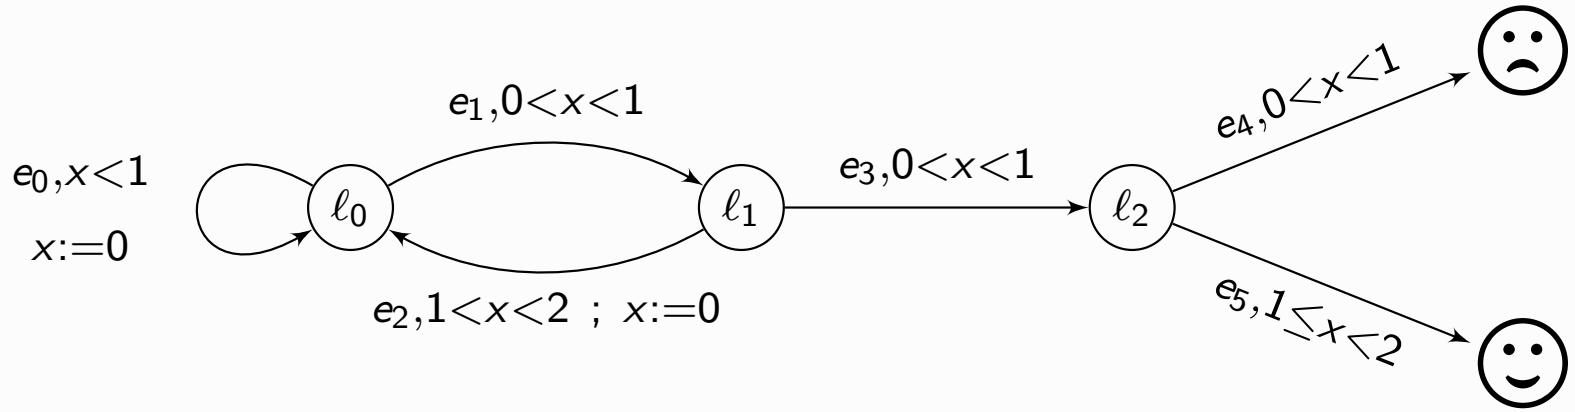

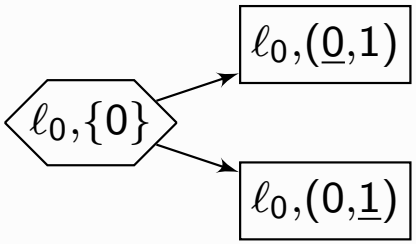

Limit corner-point region MDP

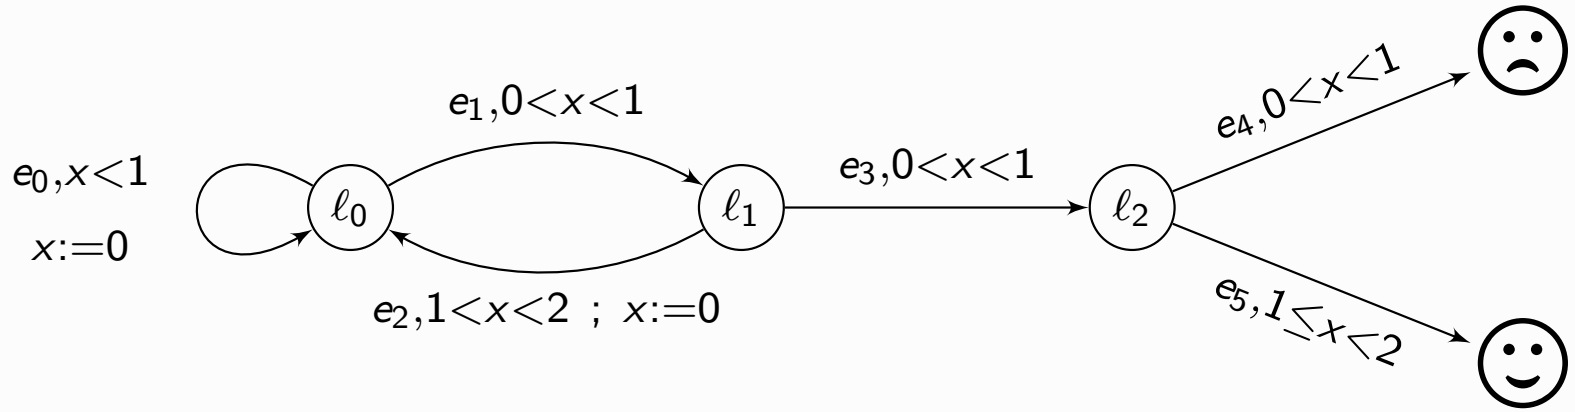

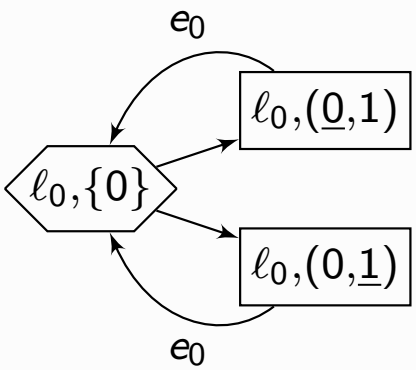

Limit corner-point region MDP

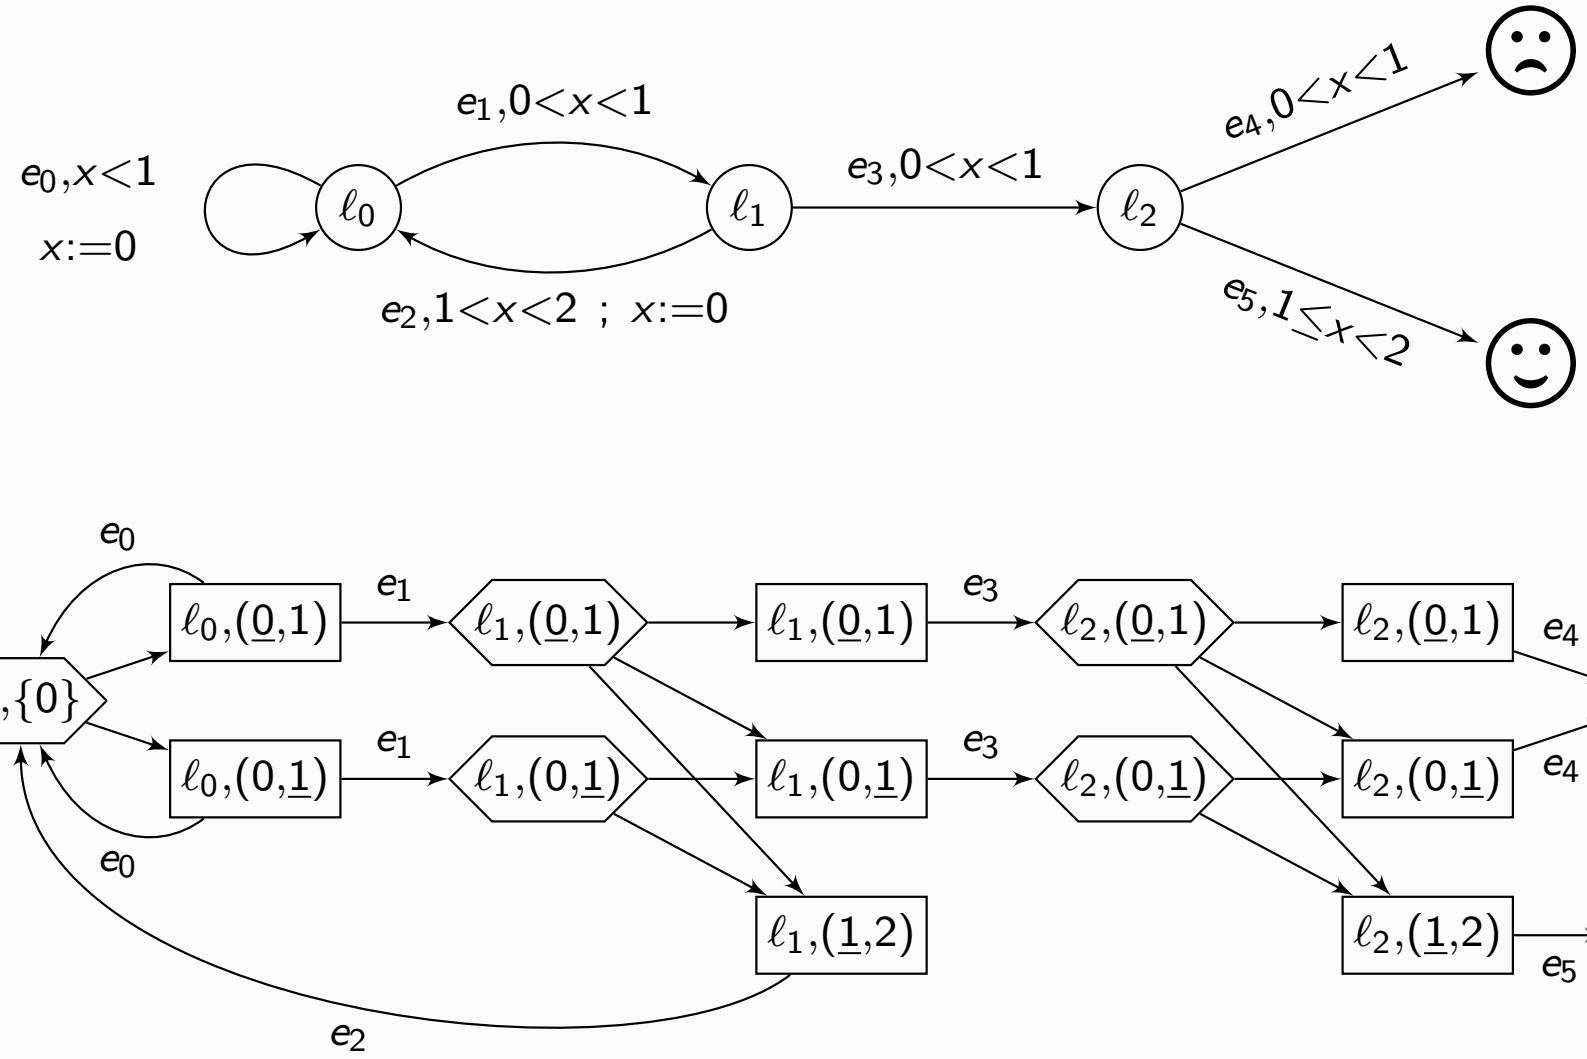

Limit corner-point region MDP

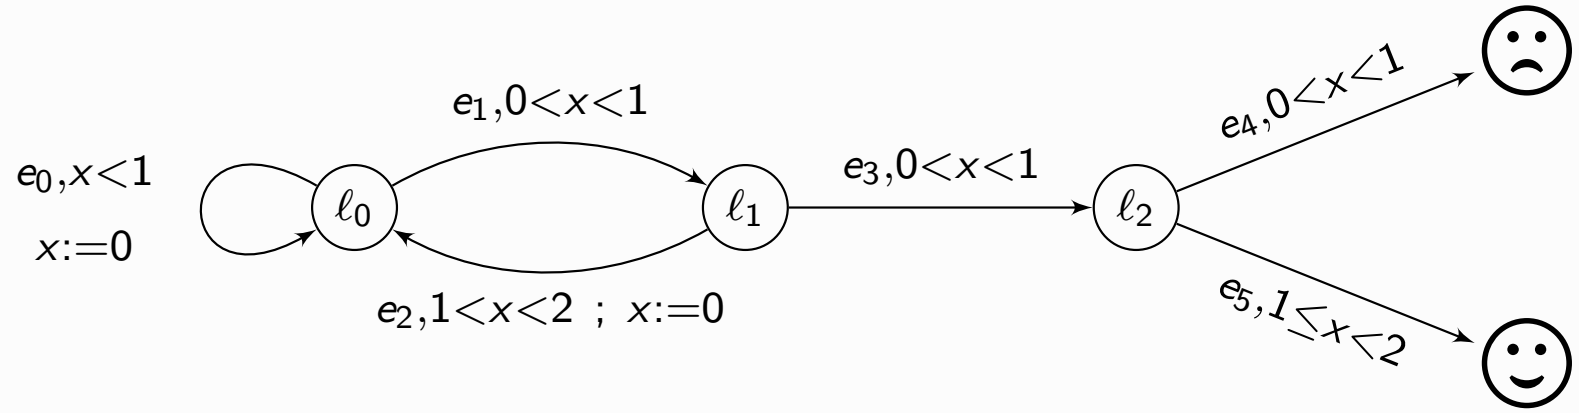

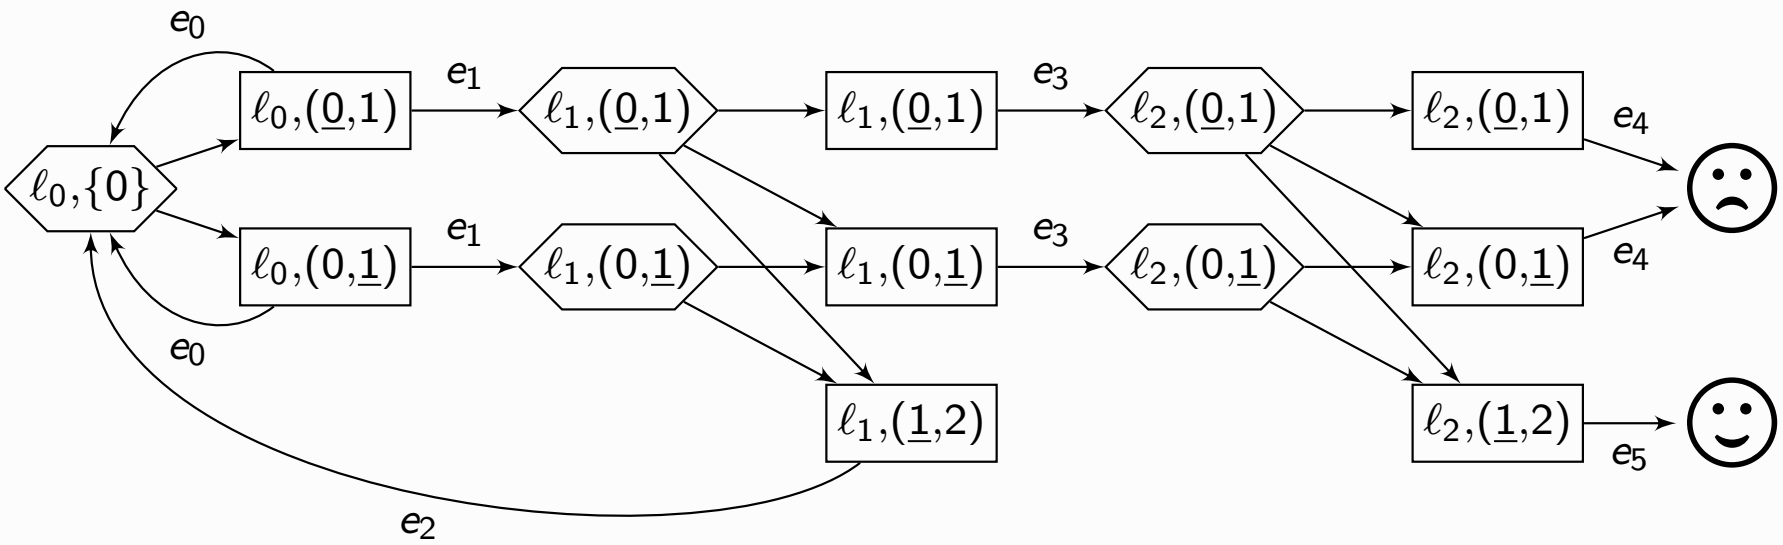

☺ is not almost-surely reachable from $\langle \ell_0, \{0\} \rangle$

Limit corner-point region MDP

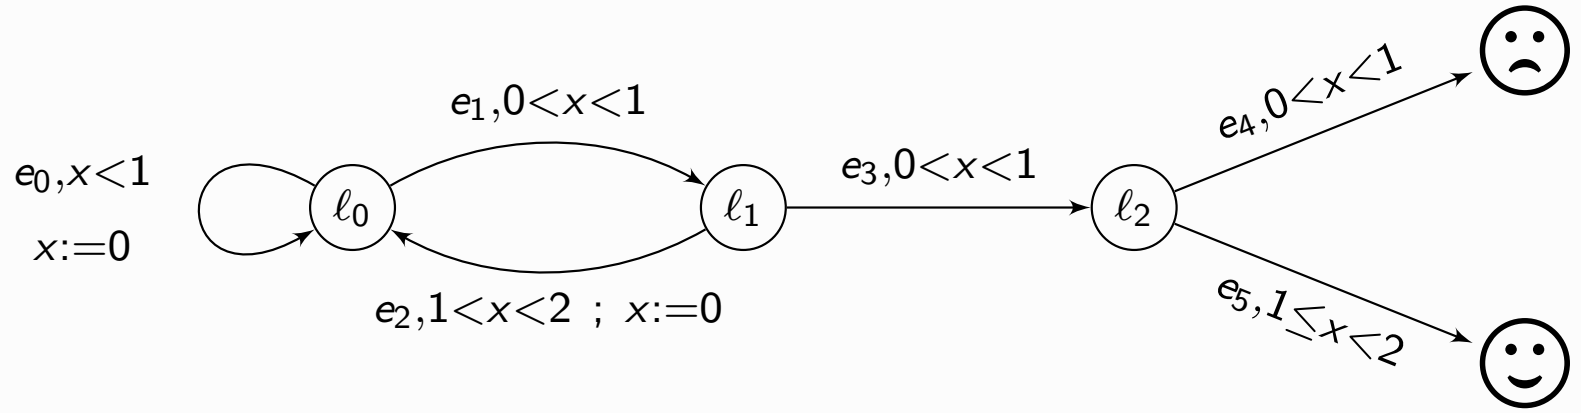

need to take into account **limit behaviours**.

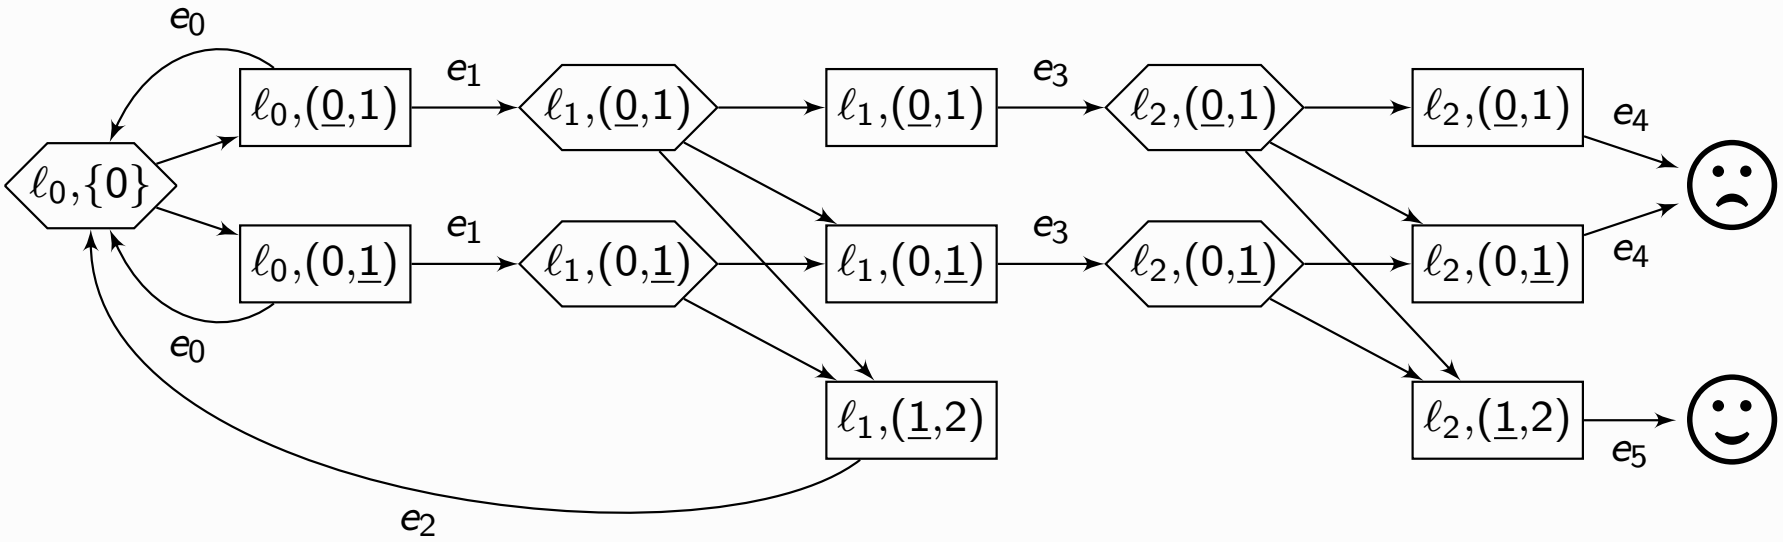

☺ is not almost-surely reachable from $\langle \ell_0, \{0\} \rangle$

Limit corner-point region MDP

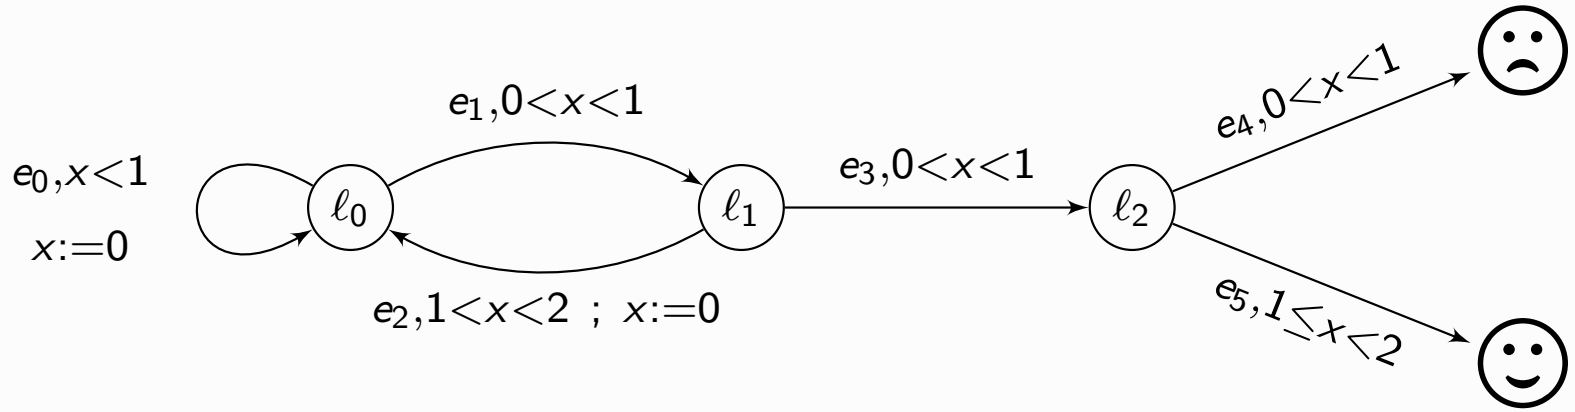

need to take into account **limit behaviours**.

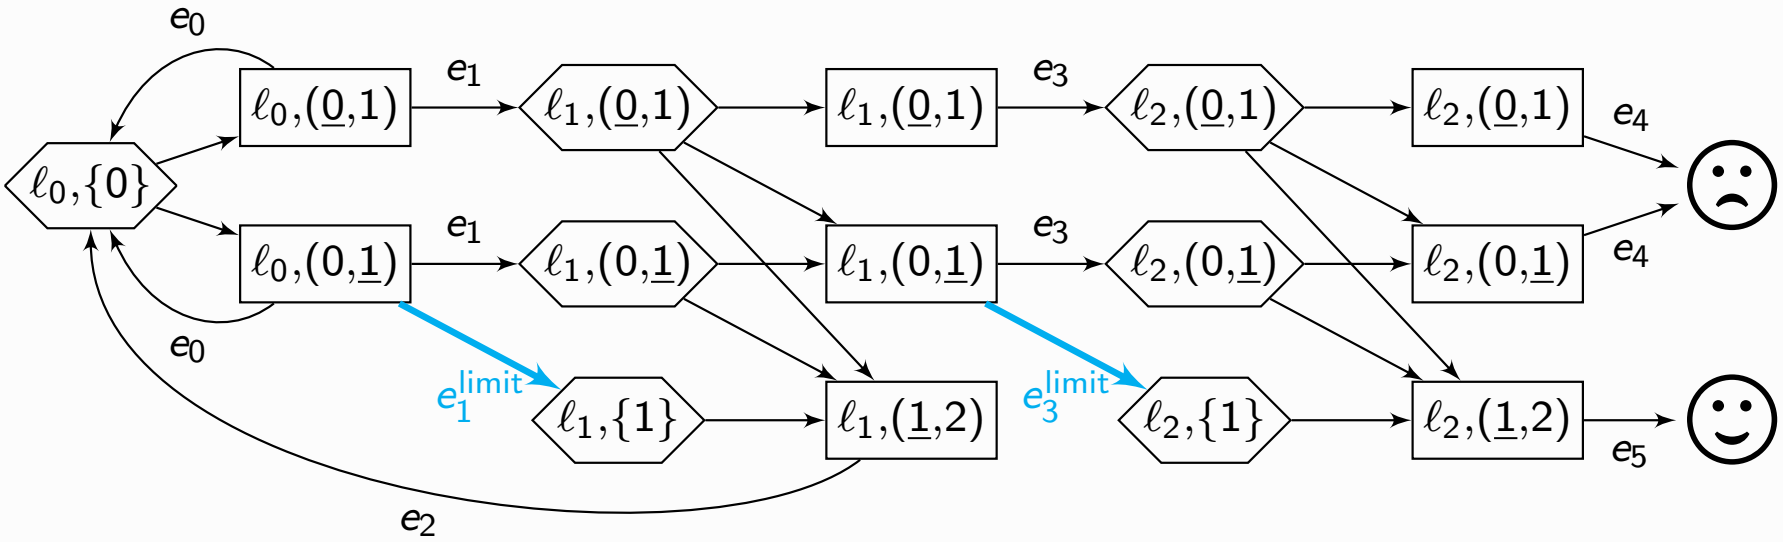

Limit corner-point region MDP

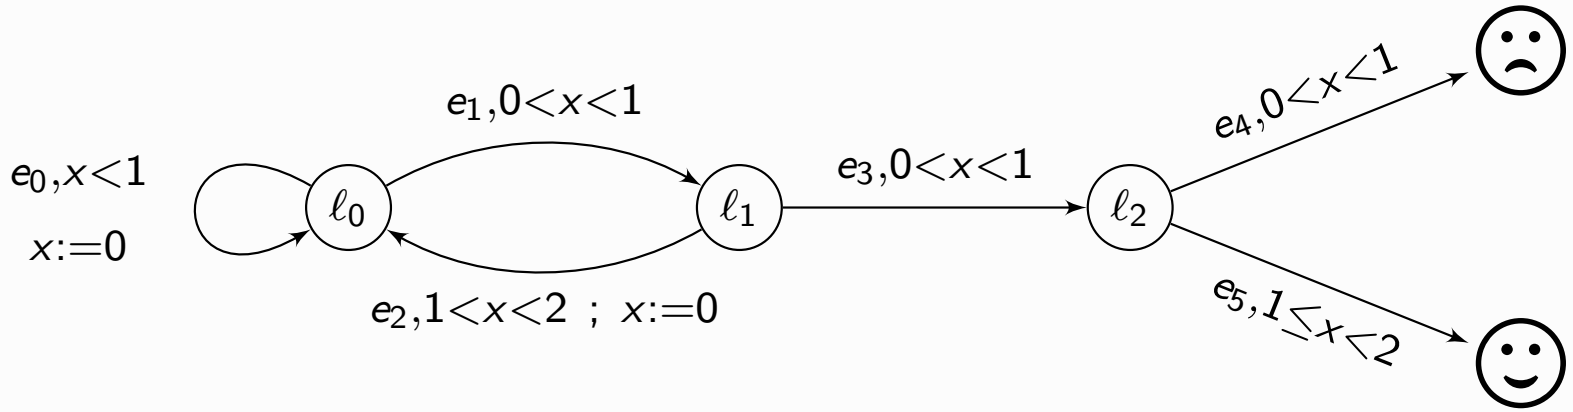

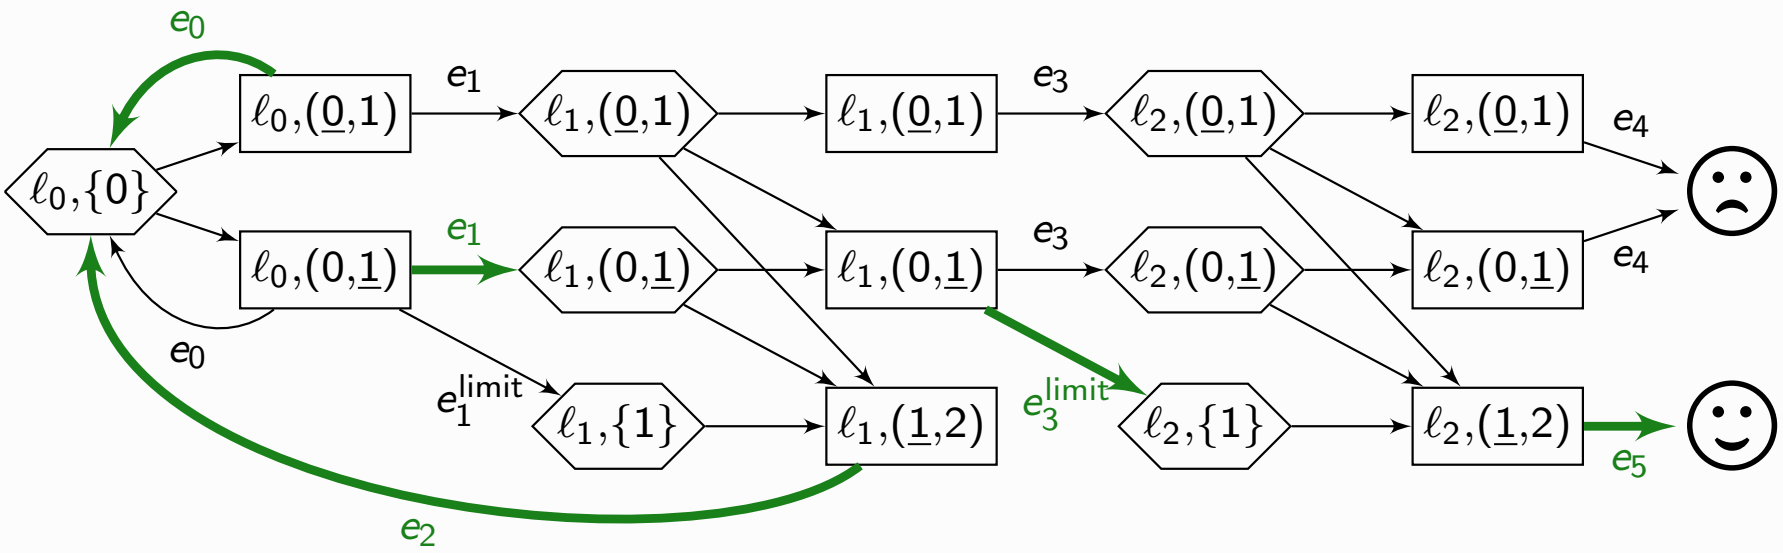

☺ is almost-surely reachable from $(\ell_0,\{0\})$

limit-sure in DSTA $\Rightarrow$ almost-sure in MDP

Proposition

If ☺ is **not almost-surely** reachable from s_0 in $\mathcal{A}_{\mathsf{cp}}$,
then ☺ is **not limit-surely** reachable from s_0 in $(\mathcal{A}, \mu)$.

Proof idea

- if $[\ell, (c, \underline{c+1})]$ is losing in MDP, then the value is uniformely bounded away from 1 for all states in $(\ell, (c, c+1))$;
- else, if $[\ell, (\underline{c}, c+1)]$ is losing in MDP, then for all states in $(\ell, (c, c+1))$ the value is bounded away from 1.

almost-sure in MDP $\Rightarrow$ limit-sure in DSTA

Proposition

If ☺ is **almost-surely** reachable from s_0 in $\mathcal{A}_{\mathsf{cp}}$,
then ☺ is **limit-surely** reachable from s_0 in $(\mathcal{A}, \mu)$.

Proof idea

- **Key**: strategies that are positional and uniform inside $(\ell, (c, c+T))$ and $(\ell, (c+T, c+1))$ suffice
- from an almost-surely winning strategy σ_{cp} in $\mathcal{A}_{\mathsf{cp}}$
  - build an abstract family of strategies σ_T in $\mathcal{A}$ such that:

$$\sigma_T(\ell, \nu) = \begin{cases} \sigma_{\mathsf{cp}}(\ell, (\underline{c}, c+1)) & \text{if } \nu \in (c, c+T) \\ \sigma_{\mathsf{cp}}(\ell, (c, \underline{c+1})) & \text{otherwise} \end{cases}$$

  - given ε, tune T to ensure probability $\geq 1 - \varepsilon$

Solving the limit corner-point MDP

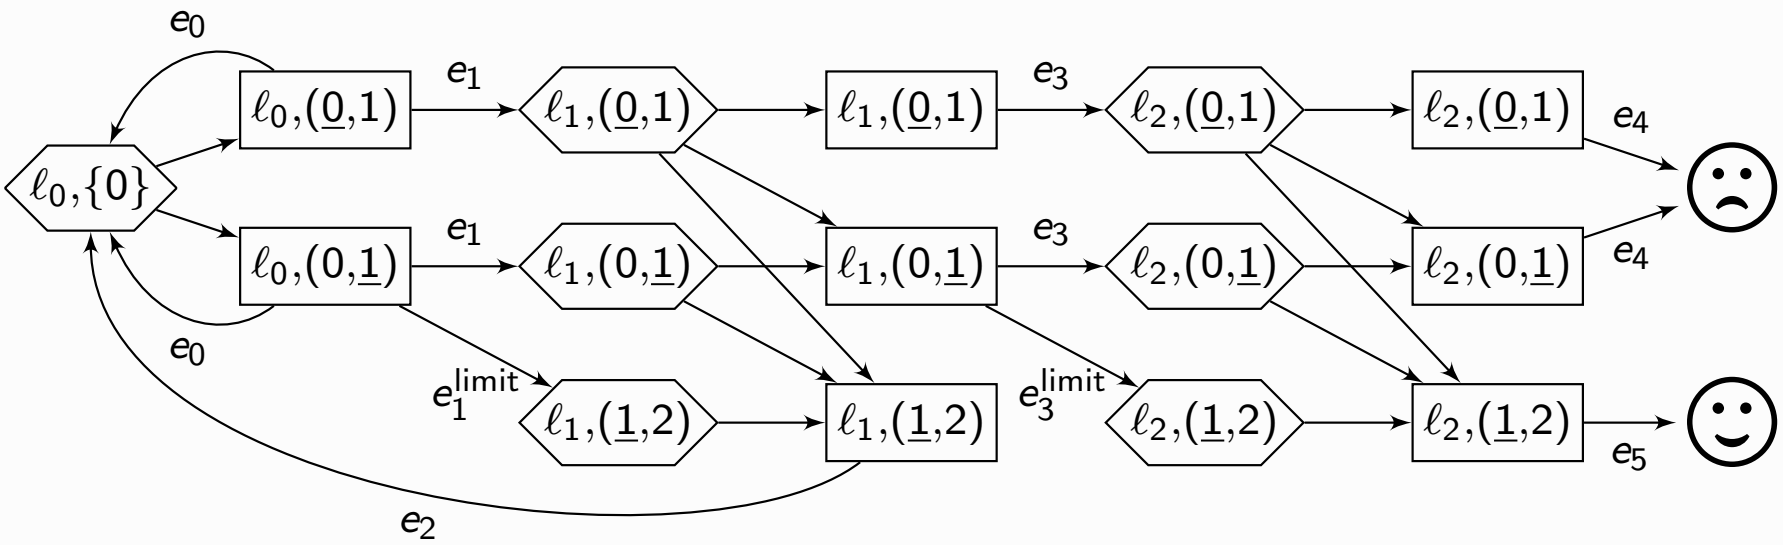

Computing winning states

Solving the limit corner-point MDP

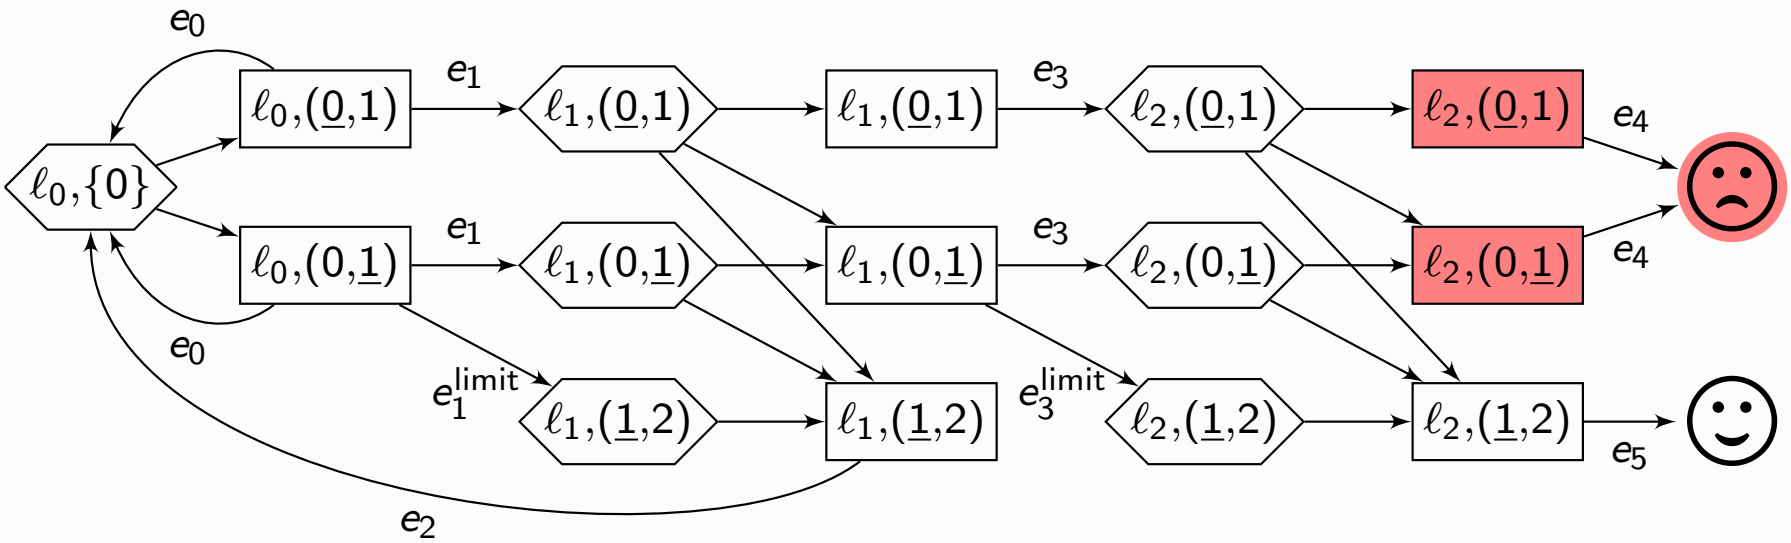

Computing winning states

- states that cannot reach ☺ are bad

Solving the limit corner-point MDP

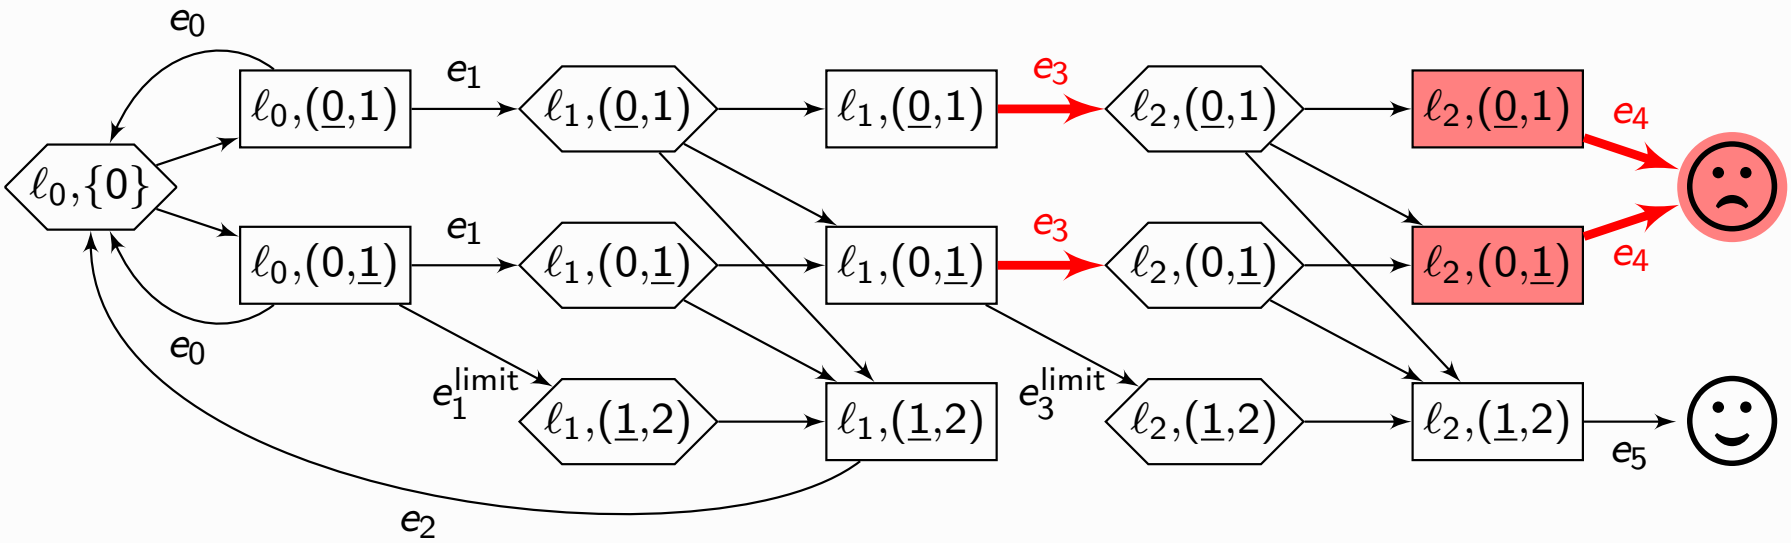

Computing winning states

- states that cannot reach ☺ are bad
- actions that lead to bad states with > 0 probability are unsafe

Solving the limit corner-point MDP

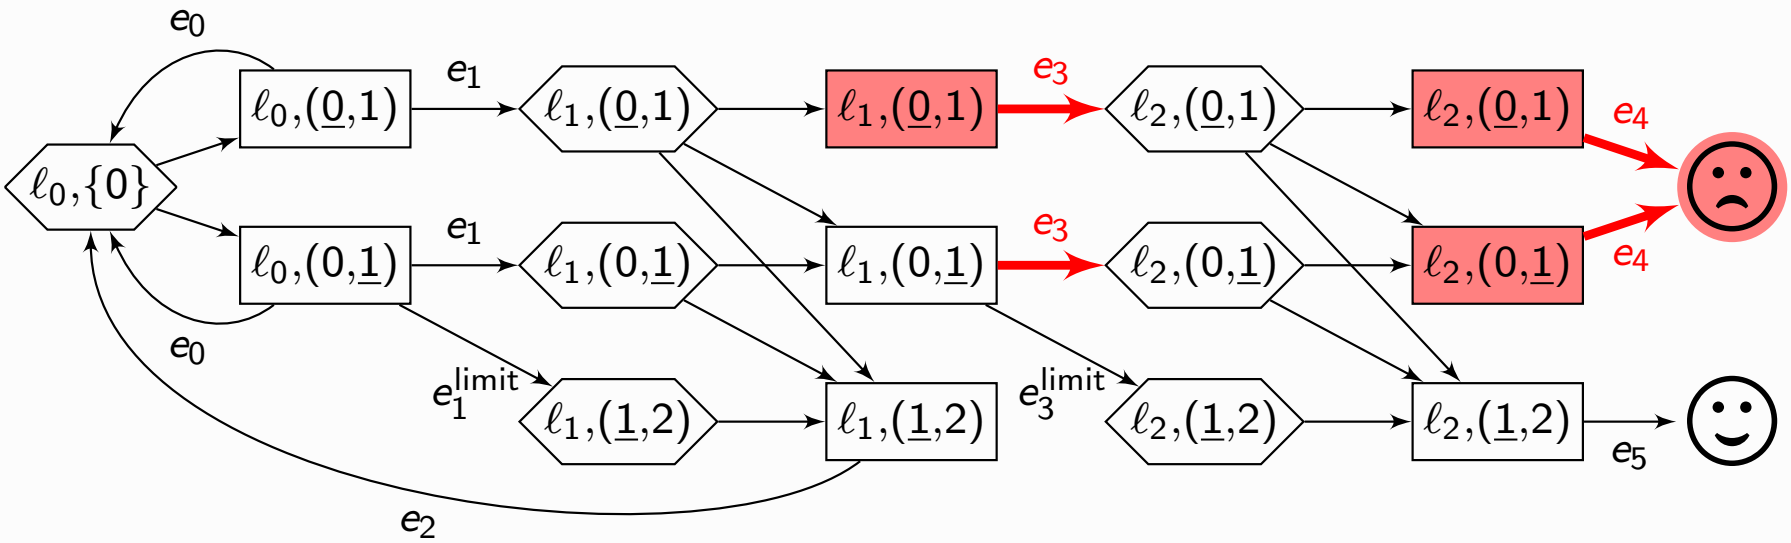

Computing winning states

- states that cannot reach ☺ are bad
- actions that lead to bad states with > 0 probability are unsafe
- states that only have unsafe actions are bad

Solving the limit corner-point MDP

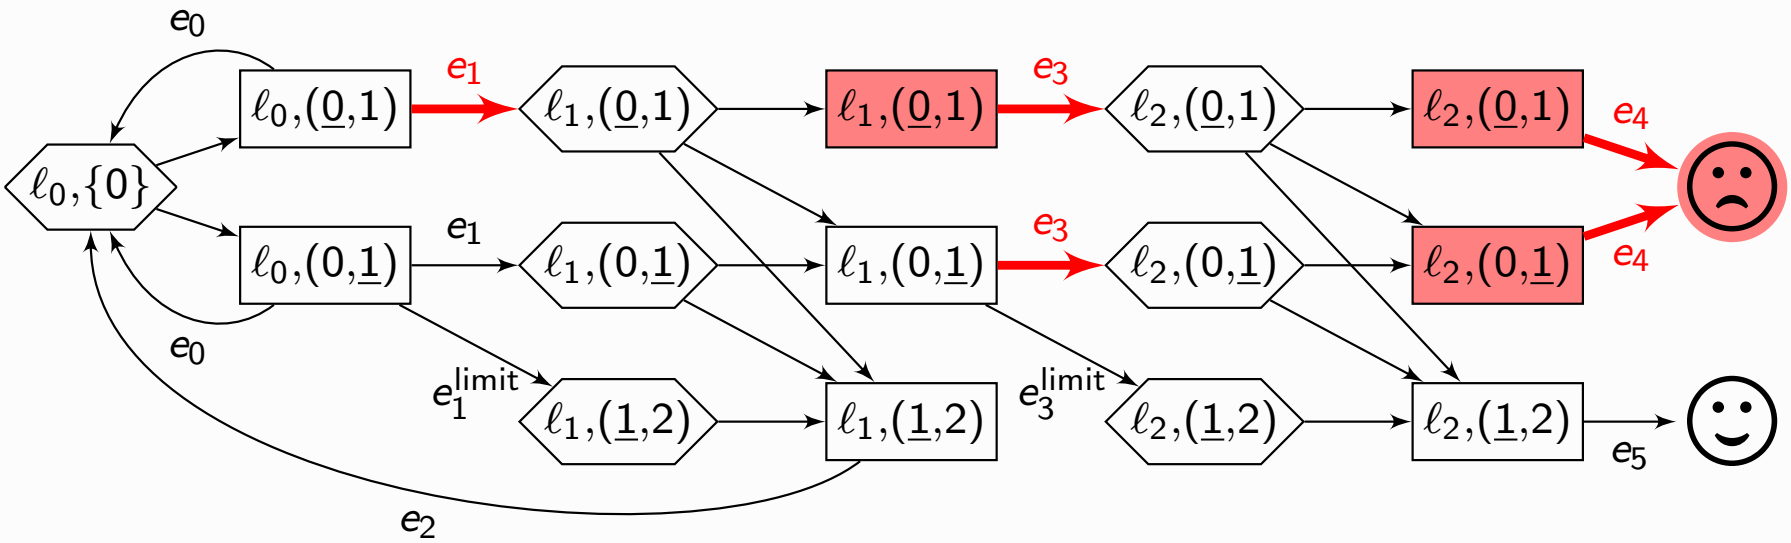

Computing winning states

- states that cannot reach ☺ are bad
- actions that lead to bad states with > 0 probability are unsafe
- states that only have unsafe actions are bad

Solving the limit corner-point MDP

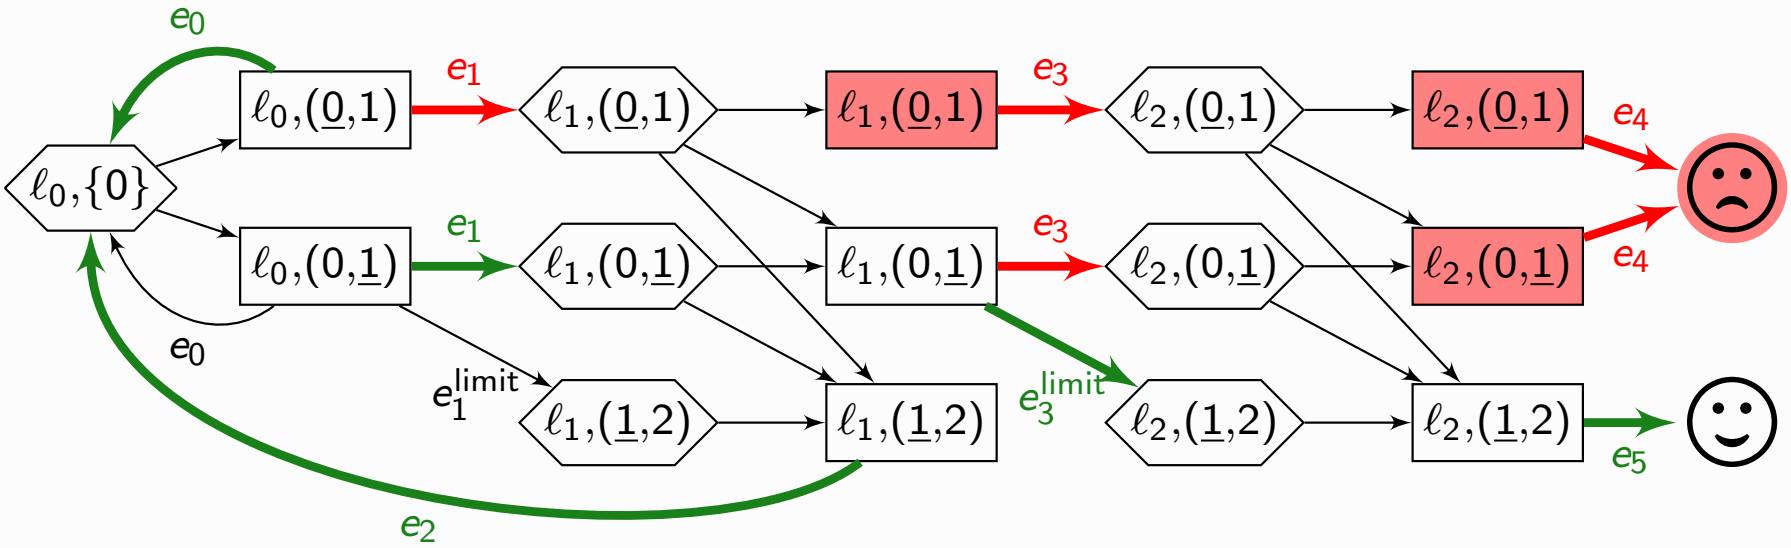

Computing winning states

- states that cannot reach ☺ are bad
- actions that lead to bad states with > 0 probability are unsafe
- states that only have unsafe actions are bad

From each winning state: safe edge towards ☺

Abstract family of strategies

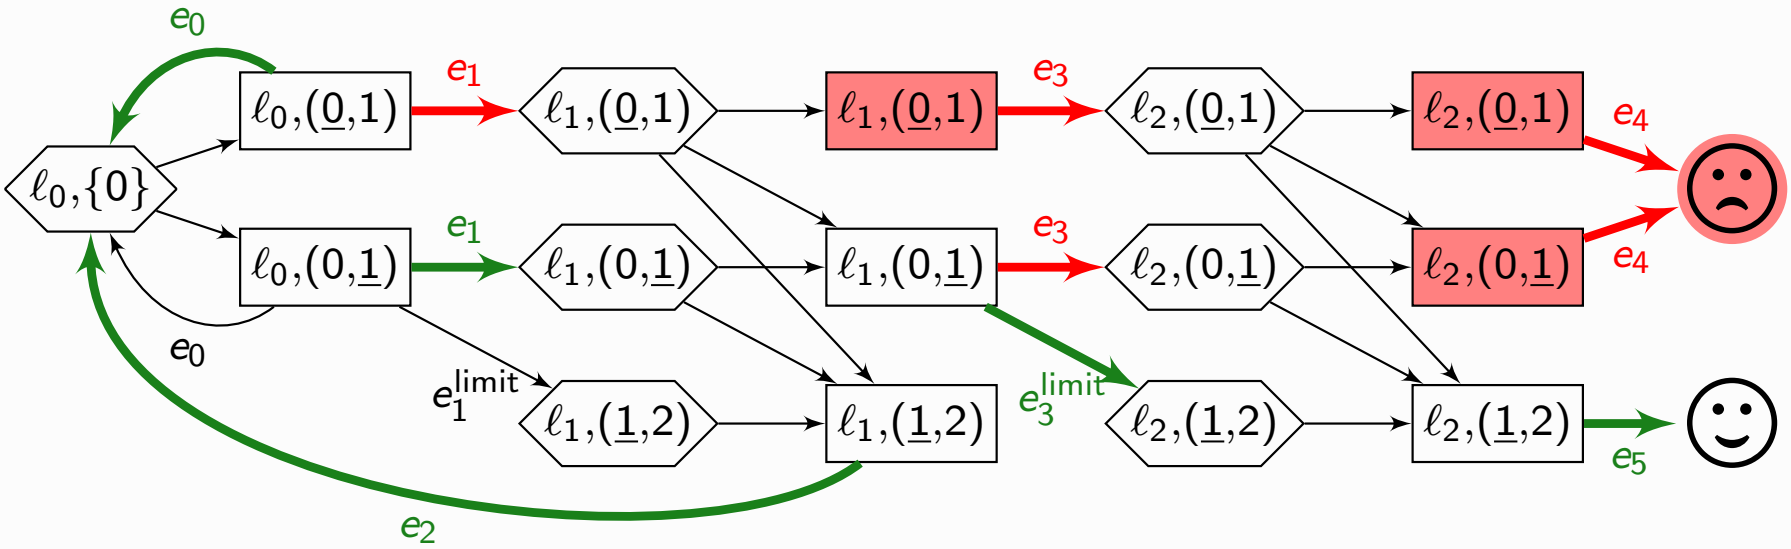

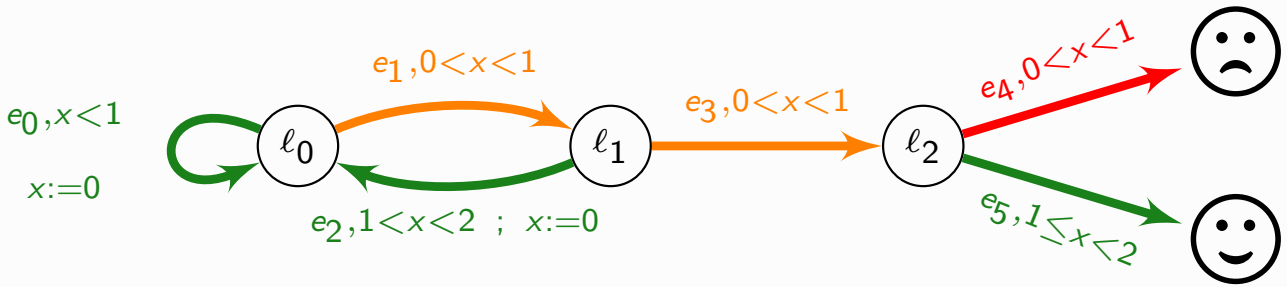

- green edges are safe
- red edges are losing
- orange edges are risky
  - $\rightsquigarrow$ chosen only when "close enough" to the right corner

▸ Details

Conclusion

Contributions

- PTIME algorithms on one-clock DSTA for
  - the almost-sure reachability problem, and
  - the limit-sure reachability problem
- non trivial ε-optimal strategies
  - not region uniform
  - cutpoint set according to "distance" to ☺

Conclusion

Contributions

- PTIME algorithms on one-clock DSTA for
  - the almost-sure reachability problem, and
  - the limit-sure reachability problem
- non trivial ε-optimal strategies
  - not region uniform
  - cutpoint set according to "distance" to ☺

Ongoing and future work

- Value 1 for other properties, and larger class of DSTA.
- Towards quantitative analysis: value approximation.

Construction of the cutpoint function T

A simple case

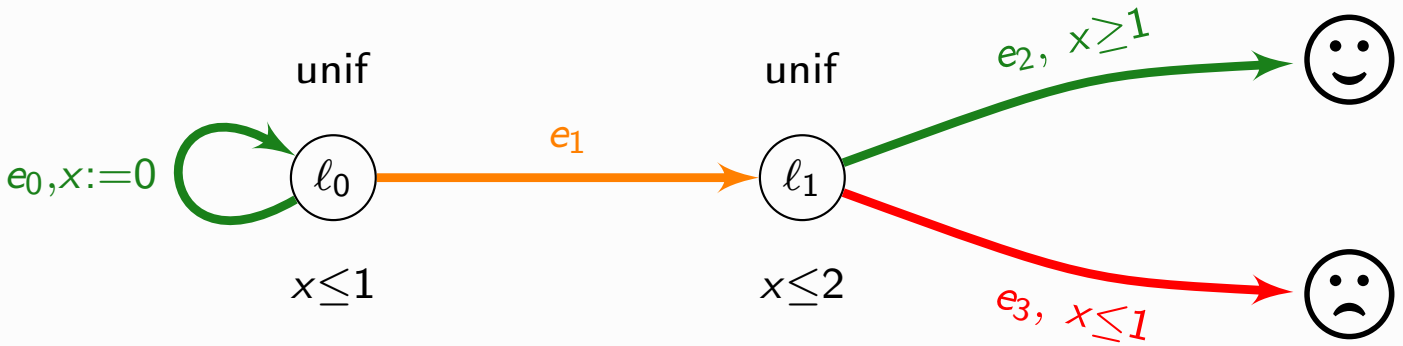

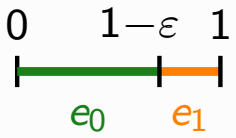

$T(\ell_0,(0,1))=1-\varepsilon$

$$\sigma_T(\ell_0,\nu) = \begin{cases} e_0 & \text{if } \nu < 1-\varepsilon \\ e_1 & \text{if } \nu \geq 1-\varepsilon \end{cases} \quad \leadsto \quad \mathbb{P}^{s_0}_{\sigma_T}\left((\mathcal{A},\mu) \models \Diamond \text{☺}\right) \geq 1-\varepsilon.$$

Construction of the cutpoint function T

A not that simple case

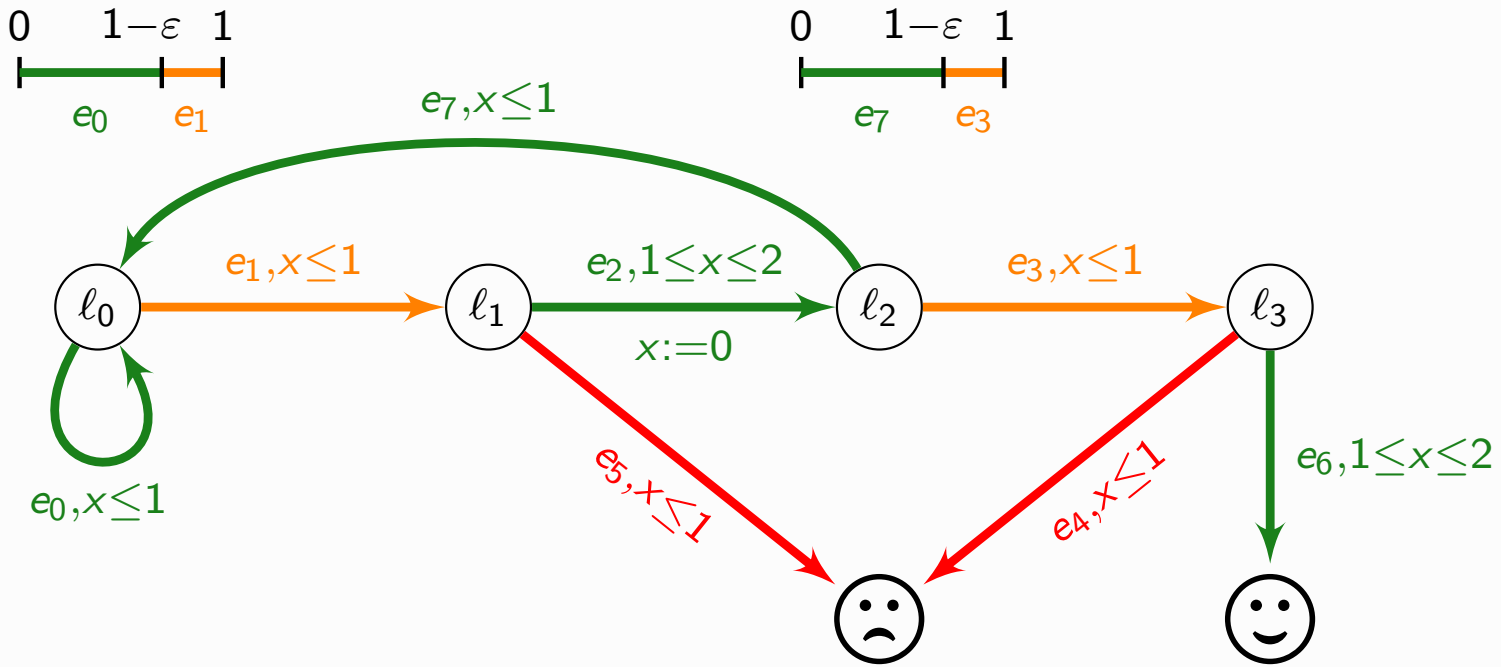

$$\mathbb{P}_{\sigma_T}([\ell_0, 0] \models \Diamond ☺) < 2/3$$

Construction of the cutpoint function T

A not that simple case

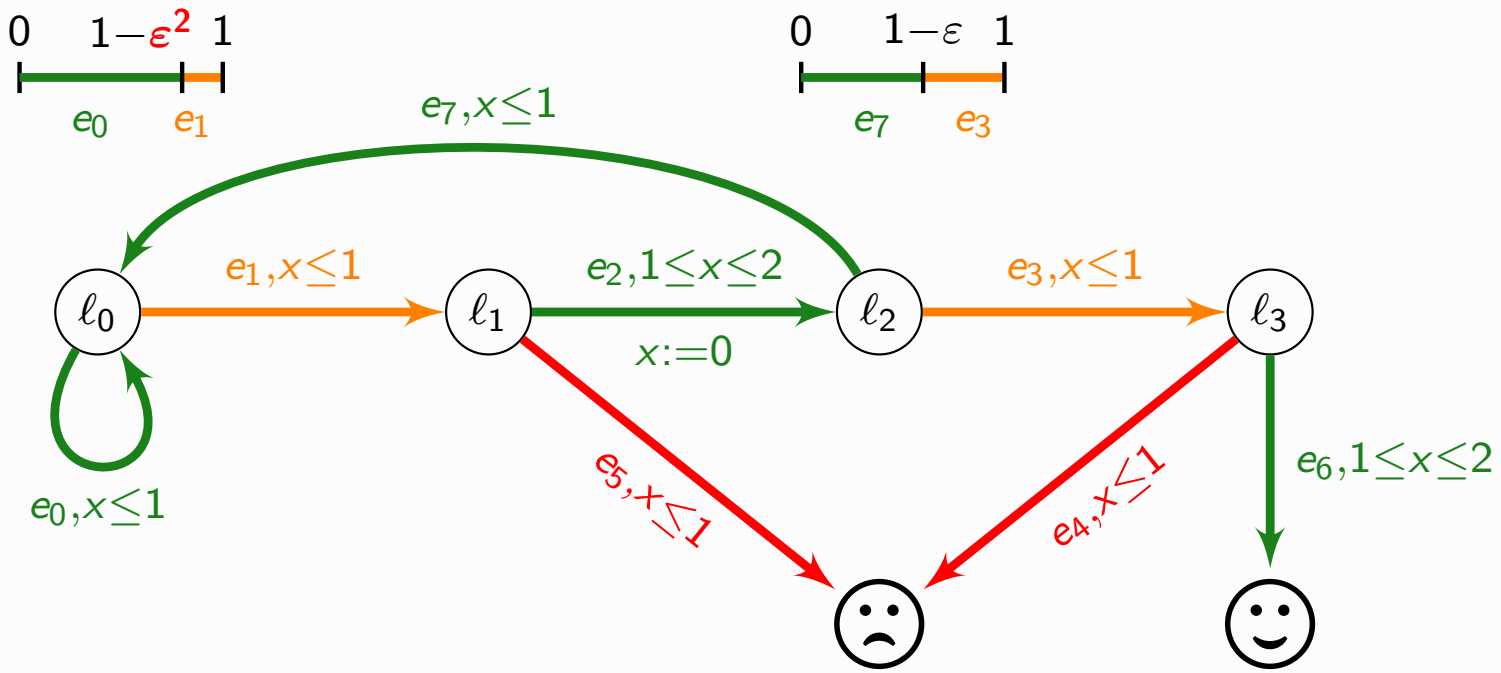

☺ is limit-surely reachable from $[\ell_0, 0]$ using involved T.

Back to main