Lecture 5
Probabilistic Systems

Nok Wongpiromsarn
UT Austin/Iowa State

Richard M. Murray
Caltech

EECI-IGSC, 10 March 2020

Outline

- Stochastic models: Markov chains, Markov decision processes
- σ-algebras
- Reachability, regular safety and ω-regular properties
- PCTL

Markov chains

A (discrete-time) Markov chain is a tuple $\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L)$ where

- S is a countable, nonempty set of states,
- $\mathbf{P} : S \times S \to [0,1]$ is the transition probability function such that for any state $s \in S$, $\sum_{s' \in S} \mathbf{P}(s, s') = 1$,
- $\iota_{init} : S \to [0,1]$ is the initial distribution such that $\sum_{s \in S} \iota_{init}(s) = 1$,
- AP is a set of atomic propositions, and
- $L : S \to 2^{AP}$ is a labeling function.

$\mathcal{M}$ is call *finite* if S and AP are finite.

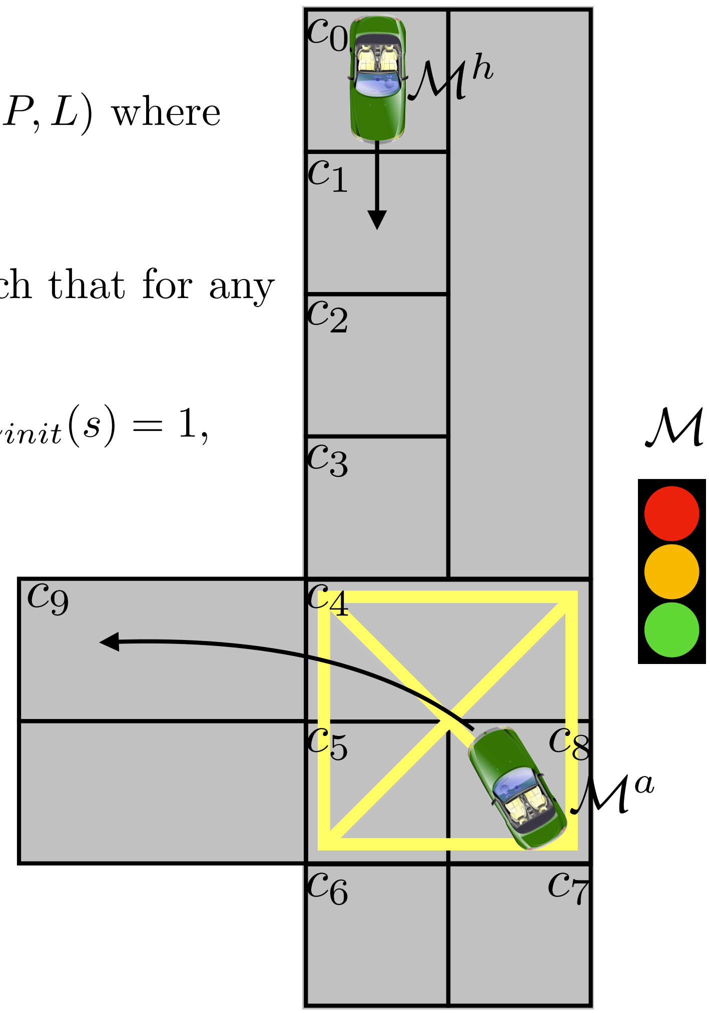

$\mathcal{M}^l$

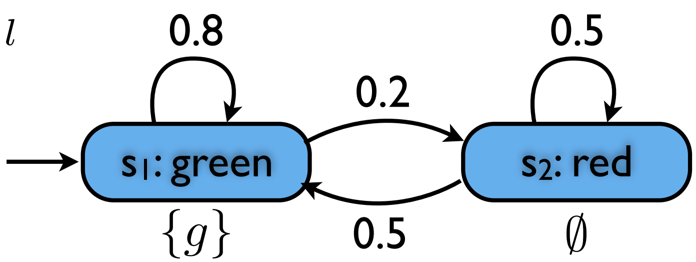

$\mathcal{M}^h$

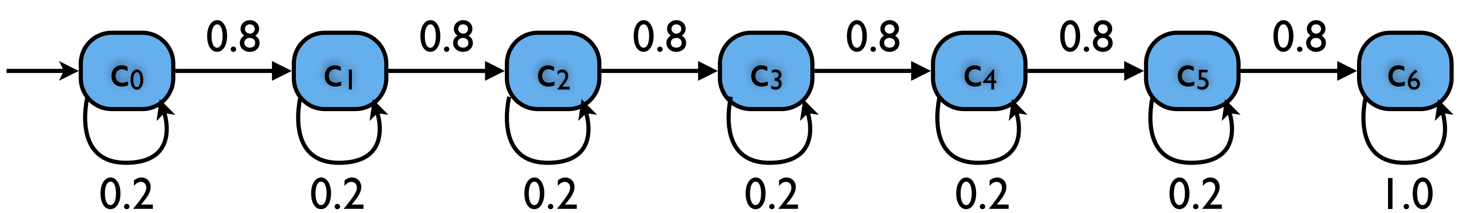

Paths of a Markov chain

Consider a Markov chain $\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L)$.
For $s \in S$,

$$Post(s) := \{s' \in S : \mathbf{P}(s, s') > 0\}$$

- A sequence of states, either finite $\pi = s_0 s_1 s_2 \ldots s_n$ or infinite $\pi = s_0 s_1 s_2 \ldots$ is a *path fragment* if $s_{i+1} \in Post(s_i), \forall i \geq 0$.
- A *path* is an infinite path fragment such that $\iota(s_0) > 0$.
- Given a path π in $\mathcal{M}$, $\inf(\pi)$ denotes the set of states that are visited infinitely often in π.
- Denote the set of paths in $\mathcal{M}$ by $Path(\mathcal{M})$
- Denote the set of finite path fragments in $\mathcal{M}$ by $Path_{fin}(\mathcal{M})$.

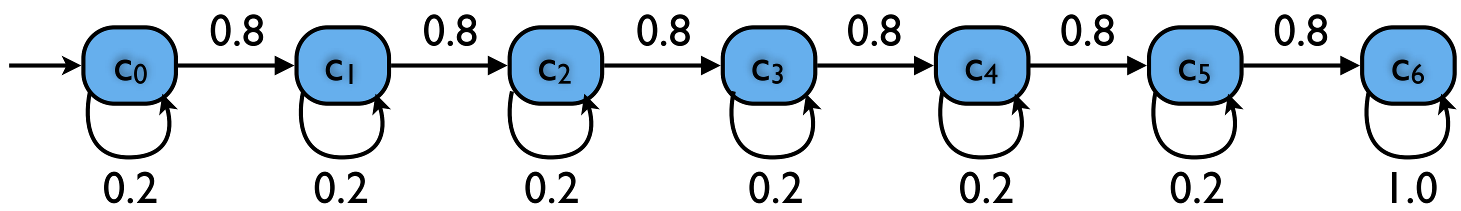

A path:

$c_0\ c_1\ c_2\ c_3\ c_4\ c_5\ c_6\ c_6 \cdots$

$c_0\ c_0\ c_0 \cdots$

Not a path:

$c_0\ c_1\ c_2\ c_3\ c_4\ c_5\ c_6$

$c_1\ c_1\ c_1 \cdots$

σ-algebras

"outcomes" "events"

Let Ω be a set. Then, $\Sigma \subseteq 2^{\Omega}$ is a σ-**algebra** if

- $\emptyset \in \Sigma$,
- $A \in \Sigma$ implies $\Omega \setminus A \in \Sigma$, i.e., Σ is closed under complementation,
- $A_1, A_2, \ldots \in \Sigma$ implies $\bigcup_{i \geq 1} A_i \in \Sigma$, i.e., Σ is closed under countable unions.

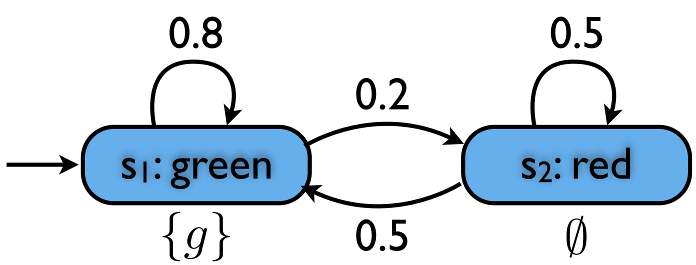

$$\Omega = \{s_1, s_2\}$$

$$\Sigma = 2^{\Omega} = \{\emptyset, \{s_1\}, \{s_2\}, \{s_1, s_2\}\}$$

Probability spaces

(Kolmogorov's) Axioms of Probability: Let Σ be a σ-algebra for some outcome space Ω. A probability measure Pr is a function from Σ to the extended real number line satisfying the following properties.

- **First Axiom:** For any $A \in \Sigma$, $Pr(A) \in \mathbb{R}$ and $Pr(A) \geq 0$
- **Second Axiom:** $Pr(\Omega) = 1$
- **Third Axiom:** If $\{A_i\}$ is a countable pairwise disjoint set with $A_i \in \Sigma$, then

$$Pr\left(\bigcup_i A_i\right) = \sum_i Pr(A_i)$$

A **probability space** is a triple (Ω, Σ, Pr).

For countable Ω, define a **distributions** on Ω as $\mu : \Omega \to [0,1]$ such that

$$\sum_{out \in \Omega} \mu(out) = 1$$

It induces a probability measure Pr on the σ-algebra $\Sigma = 2^{\Omega}$ defined as

$$Pr(A) = \sum_{out \in A} \mu(out), \forall A \in \Sigma$$

Probability measures of a Markov chain

Consider a Markov chain $\mathcal{M}$.

- $\Omega = Path(\mathcal{M})$ plays the role of the outcomes.
- Define a **cylinder set** of $\hat{\pi} = s_0 \ldots s_n \in Path_{fin}(\mathcal{M})$ as

$$Cyl(\hat{\pi}) = \{\pi \in Path(\mathcal{M}) \mid \hat{\pi} \in pref(\pi)\}$$

- The σ-algebra associated with $\mathcal{M}$ is the smallest σ-algebra that contains all $Cyl(\hat{\pi}), \hat{\pi} \in Path_{fin}(\mathcal{M})$
- There exists a unique probability measure $Pr^{\mathcal{M}}$ such that

$$Pr^{\mathcal{M}}(Cyl(s_0 \ldots s_n)) = \iota_{init}(s_0) \prod_{0 \leq i < n} \mathbf{P}(s_i, s_{i+1})$$

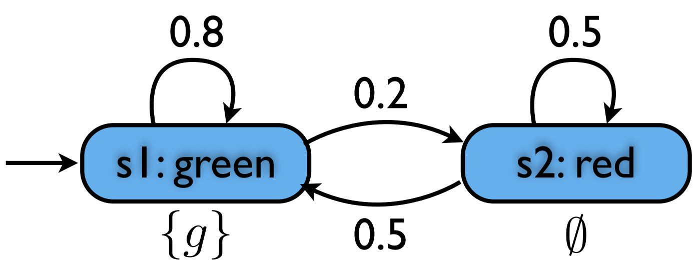

Repeat finitely many times

Repeat infinitely many times

$$\hat{\pi} = s_1 s_1 s_2$$

$$Cyl(\hat{\pi}) = \{s_1 s_1 s_2 (s_1^* s_2^*)^\omega\}$$

$$Pr(Cyl(\hat{\pi})) = 1 \cdot \mathbf{P}(s_1, s_1) \cdot \mathbf{P}(s_1, s_2) = 1 \cdot 0.8 \cdot 0.2$$

Probability of satisfying an LTL formula

Consider an LTL formula φ over AP and an MC $\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L)$. The probability of $\mathcal{M}$ satisfying φ is given by

$$Pr^{\mathcal{M}}(\varphi) = Pr^{\mathcal{M}}(\{\pi \in Path(\mathcal{M}) \mid \pi \models \varphi\})$$

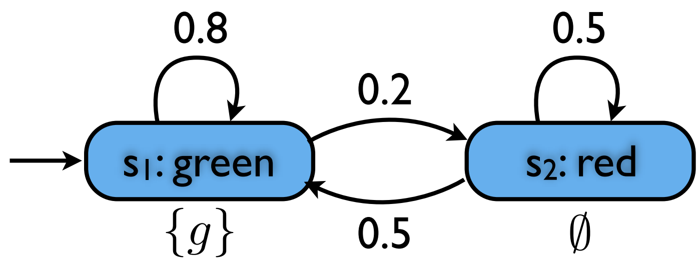

$$\begin{aligned}
Pr^{\mathcal{M}}(\Diamond \neg g) &= \sum_{s_0 \ldots s_n \in Path_{fin}(\mathcal{M}) \cap \{s_1^* s_2\}} Pr^{\mathcal{M}}(Cyl(s_0 \ldots s_n)) \\
&= \sum_{s_0 \ldots s_n \in Path_{fin}(\mathcal{M}) \cap \{s_1^* s_2\}} \iota_{init}(s_0) \prod_{0 \le i < n} \mathbf{P}(s_i, s_{i+1}) \\
&= \sum_{i=0}^{\infty} (0.8)^i \cdot 0.2 \\
&= \frac{0.2}{1 - 0.8} \\
&= 1
\end{aligned}$$

$\Diamond \neg g$ holds **almost surely**

Probabilistic Computation Tree Logic (PCTL)

- Recall that CTL includes the path quantifiers $\exists$ and $\forall$
- PCTL replaces the path quantifiers with the probabilistic operator $\mathbb{P}_J(\varphi)$ where $J \subseteq [0,1]$ is an interval with rational bounds

$$s \models \mathbb{P}_J(\varphi) \text{ iff } Pr(s \models \varphi) \in J$$

$$Sat\big(\mathbb{P}_J(\varphi)\big) = \{s \in S \mid Pr(s \models \varphi) \in J\}$$

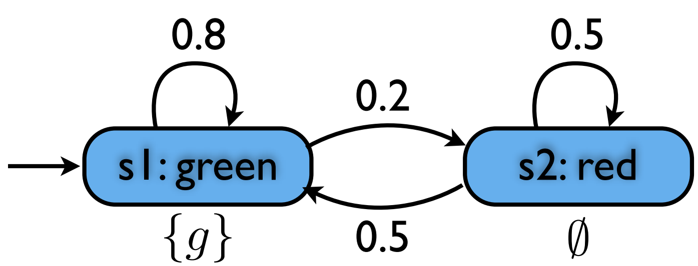

Recall $Pr(\Diamond \neg g) = 1$

$$Sat\big(\mathbb{P}_{[1,1]}(\Diamond \neg g)\big) = \{s1, s2\}$$

Reachability property

Consider a Markov chain $\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L)$. For each $s \in S$, define a Markov chain $\mathcal{M}_s = (S, \mathbf{P}, \iota_{init,s}, AP, L)$ where

$$\iota_{init,s}(t) = \begin{cases} 1 & \text{if } t = s \\ 0 & \text{otherwise} \end{cases}$$

The probability for φ to hold in a state s is given by

$$Pr^{\mathcal{M}}(s \models \varphi) = Pr^{\mathcal{M}_s}\Big(\{\pi \in Path(\mathcal{M}_s) \mid \pi \models \varphi\}\Big)$$

For each $s \in S$, define $x_s = Pr^{\mathcal{M}}(s \models \Diamond B)$ where $B \subseteq S$.

- If B is not reachable from s, $x_s = 0$.
- $x_s = 1$ for all $s \in B$.
- Define $\tilde{S} = \{s \in S \setminus B \mid B \text{ is reachable from } s\}$. For any $s \in \tilde{S}$,

$$x_s = \sum_{t \in S \setminus B} \mathbf{P}(s,t) \cdot x_t + \sum_{u \in B} \mathbf{P}(s,u)$$

reaching a state $t \in S \setminus B$, from which B is reached

reaching B within one step

Verifying reachability property

Recall that $\tilde{S} = \{s \in S \setminus B \mid B \text{ is reachable from } s\}$. For any $s \in \tilde{S}$,

$$x_s = \sum_{t \in S \setminus B} \mathbf{P}(s,t) \cdot x_t + \sum_{u \in B} \mathbf{P}(s,u)$$

Define $\mathbf{x} = (x_s)_{s \in \tilde{S}}$. In matrix form, define $\mathbf{x} = \left(x_{s \in \tilde{S}}\right)$.

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{b}$$

Transition probability matrix

Compute $\tilde{S}$ — Graph search, e.g., backward DFS or BFS

$\Downarrow$

Construct $\mathbf{A}$ and $\mathbf{b}$ — Transition probability $\mathbf{P}$

$\Downarrow$

Solve $\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{b}$ — Solve $\mathbf{x} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{b}$ if the inverse exists. Otherwise, compute the least fixed point of

$$\Gamma(\mathbf{y}) = \mathbf{A}\mathbf{y} + \mathbf{b}$$

Verifying regular safety properties

states transitions initial states

Recall NFA $\mathcal{A} = (Q,\ \Sigma,\ \delta,\ Q_0,\ F)$.

A *deterministic finite automaton (DFA)* is an NFA with

$$\begin{aligned} |Q_0| &\leq 1 \\ |\delta(q, A)| &\leq 1, \forall q \in Q, A \in \Sigma \end{aligned}$$

For any NFA, one can construct an equivalent DFA through *powerset construction*.

Consider a regular safety property P_{safe}. Let $\mathcal{A}$ be a DFA for the bad prefixes of P_{safe}.

$$Pr^{\mathcal{M}}(P_{safe}) = 1 - \sum_{s \in S} \iota_{init}(s) Pr(s \models \mathcal{A})$$

$$Pr^{\mathcal{M}_s}\left(\{\pi \in Path(\mathcal{M}_s) \mid pref(trace(\pi)) \cap \mathcal{L}(\mathcal{A}) \neq \emptyset\}\right)$$

Product Markov Chain

Markov chain

DFA

$$\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L) \quad \otimes \quad \mathcal{A} = (Q, 2^{AP}, \delta, q_0, F)$$

$$\shortparallel$$

$$(S', \mathbf{P}', \iota'_{init}, AP', L')$$

Product Markov chain

- $S' = S \times Q$
- $\mathbf{P}'(\langle s, q\rangle, \langle s', q'\rangle) = \begin{cases} \mathbf{P}(s, s') & \text{if } q \xrightarrow{L(s')} q' \\ 0 & \text{otherwise} \end{cases}$
- $\iota'_{init}(\langle s, q\rangle) = \begin{cases} \iota_{init}(s) & \text{if } q_0 \xrightarrow{L(s)} q \\ 0 & \text{otherwise} \end{cases}$
- $AP' = \{accept\}$
- $L'(\langle s, q\rangle) = \begin{cases} \{accept\} & \text{if } q \in F \\ \emptyset & \text{otherwise} \end{cases}$

For any path fragment $s_0 s_1 s_2 \ldots$ in $\mathscr{M}$, there exists a **unique** run $q_0 q_1 q_2 \ldots$ in $\mathscr{A}$ for $L(s_0)L(s_1)L(s_2)\ldots$ and

$\langle s_0, q_1\rangle\langle s_1, q_2\rangle\langle s_2, q_3\rangle\ldots$

is the corresponding unique path in $\mathscr{M} \otimes \mathscr{A}$.

$$\boxed{Pr^{\mathcal{M}}(s \models P_{safe}) = 1 - Pr^{\mathcal{M}\otimes\mathcal{A}}(\langle s, \delta(q_0, L(s))\rangle \models \Diamond accept)}$$

Deterministic Rabin automaton (DRA)

A *deterministic Buchi automaton (DBA)* is an NBA $\mathcal{A} = (Q, \Sigma, \delta, Q_0, F)$ with $|Q_0| \leq 1$ and $|\delta(q, A)| \leq 1$ for all $q \in Q, A \in \Sigma$.

A *deterministic Rabin automaton (DRA)* $\mathcal{A} = (Q, \Sigma, \delta, Q_0, Acc)$ has the same components as a DBA but with acceptance condition given by a set of **pairs** of states

$$Acc = \{(L_i, K_i) \mid L_i, K_i \subseteq Q\}$$

A run is *accepting* if it satisfy the LTL formula

$$\bigvee_i \left(\Diamond\Box\neg L_i \wedge \Box\Diamond K_i \right)$$

The states in L_i are visited **finitely** often

The states in K_i are visited **infinitely** often

- A DBA is a DRA with $Acc = \{(\varnothing, F)\}$
- Some ω-regular properties (e.g., $\Diamond\Box a$) cannot be expressed by a DBA
- The class of languages accepted by DRAs agrees with the class of ω-regular languages

Strongly connected components of Markov chains

$\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L)$

- $T \subseteq S$ is *strongly connected* if for each pair $s, t \in T$, there exists a finite path fragment $s_0 s_1 \ldots s_n$ such that $s_0 = s$, $s_n = t$ and $s_i \in T, \forall i$.
- A *strongly connected component (SCC)* of $\mathcal{M}$ is a subset $T \subseteq S$ such that T is strongly connected and there does not exists $T' \supset T$ that is strongly connected.
- A *bottom SCC (BSCC)* is an SCC T from which no state outside T is reachable, i.e., $\sum_{t \in T} \mathbf{P}(s,t) = 1$ for all $s \in T$.
- Denote the set of all BSCCs of $\mathcal{M}$ by $BSCC(\mathcal{M})$.

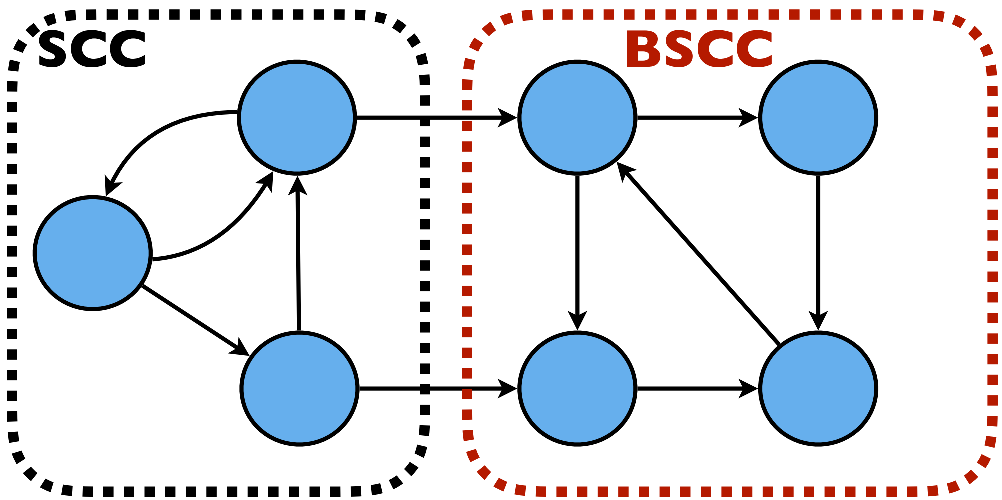

DRA-based analysis of Markov chains

Markov chain

DRA

$$\mathcal{M} = (S, \mathbf{P}, \iota_{init}, AP, L) \quad \otimes \quad \mathcal{A} = (Q, 2^{AP}, \delta, q_0, Acc)$$

$$= (S', \mathbf{P}', \iota'_{init}, AP', L')$$

Product Markov chain

- S', $\mathbf{P}'$, ι'_{init} are defined as in the product of Markov chain and DFA
- $AP' = \{L_1, \ldots, L_k, K_1, \ldots, K_k\}$ where $\{(L_1, K_1), \ldots, (L_k, K_k)\} = Acc$
- $L'(\langle s, q\rangle) = \{H \in AP' \mid q \in H\}$
- A BSCC T in $\mathcal{M} \otimes \mathcal{A}$ is *accepting* if there exists $i \in \{1, \ldots, k\}$ such that

$$\boxed{T \cap (S \times L_i) = \emptyset \quad \text{and} \quad T \cap (S \times K_i) \neq \emptyset}$$

Once T is reached, the acceptance criterion for $\mathcal{A}$ is satisfied almost surely

The union of all accepting BSCCs in $\mathcal{M} \otimes \mathcal{A}$

$$\boxed{Pr^{\mathcal{M}}(s \models \mathcal{A}) = Pr^{\mathcal{M} \otimes \mathcal{A}}\big(\langle s, \delta(q_0, L(s))\rangle \models \Diamond U\big)}$$

Verifying ω-regular property

$$Pr^{\mathcal{M}}(s \models \mathcal{A}) = Pr^{\mathcal{M}\otimes\mathcal{A}}\big(\langle s, \delta(q_0, L(s))\rangle \models \Diamond U\big)$$

Compute BSCCs of $\mathscr{M} \otimes \mathscr{A}$

Graph analysis, e.g., Tarjan's algorithm

Compute U

Check each BSCC agains all pairs $(L_i, K_i) \in Acc$

Compute reachability probabilities

Solve a set of linear equations

Markov decision processes (MDPs)

A Markov decision process is a tuple $\mathcal{M} = (S, Act, \mathbf{P}, \iota_{init}, AP, L)$ where

- S, ι_{init}, AP and L are defined as in the definition of a Markov chain,
- Act is a set of actions, and
- $\mathbf{P} : S \times Act \times S \to [0, 1]$ is the transition probability function such that for any state $s \in S$ and action $\alpha \in Act$, $\sum_{s' \in S} \mathbf{P}(s, \alpha, s') \in \{0, 1\}$.

An MDP is *finite* if S, Act and AP are finite.

An action α is *enabled* in state s iff $\sum_{s' \in S} \mathbf{P}(s, \alpha, s') = 1$.

$Act(s) = \{\alpha \in Act \mid \alpha \text{ is enabled in } s\} \neq \emptyset$

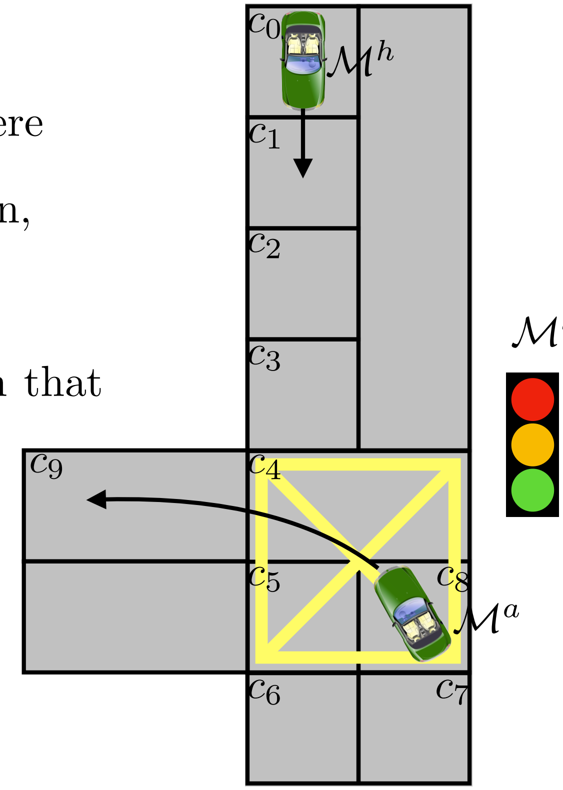

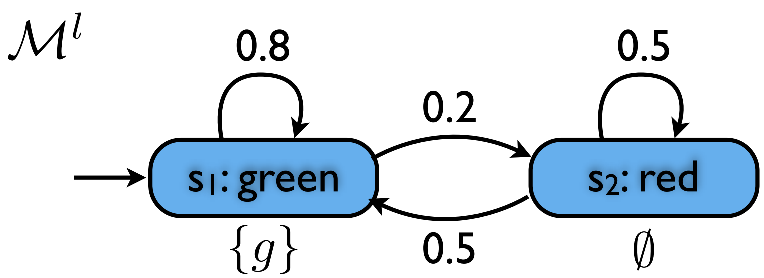

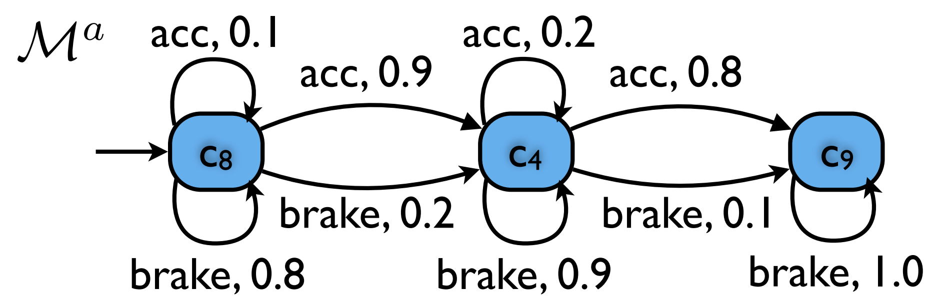

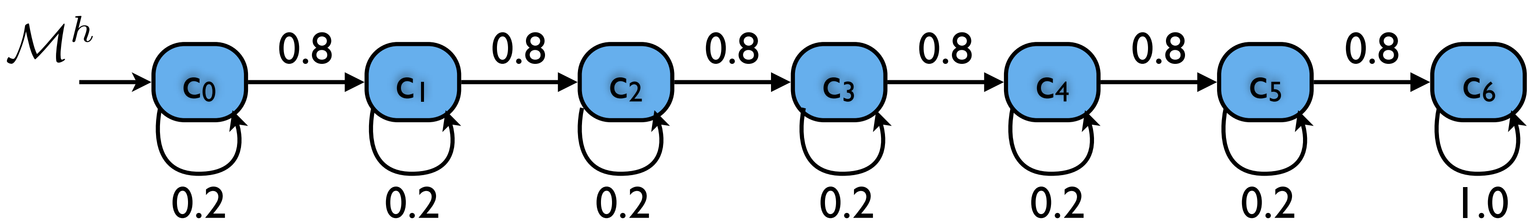

Parallel composition of MDP and MC

$\mathcal{M}^a = (S^a, Act, \mathbf{P}^a, \iota^a_{init}, AP^a, L^a)$

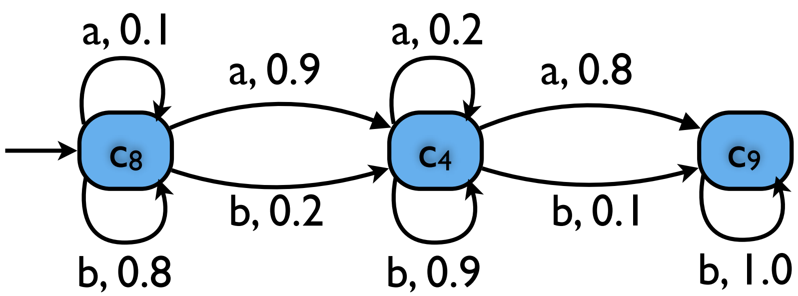

$\mathcal{M}^l = (S^l, \mathbf{P}^l, \iota^l_{init}, AP^l, L^l)$

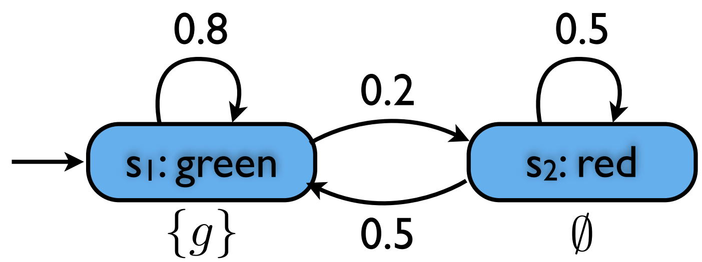

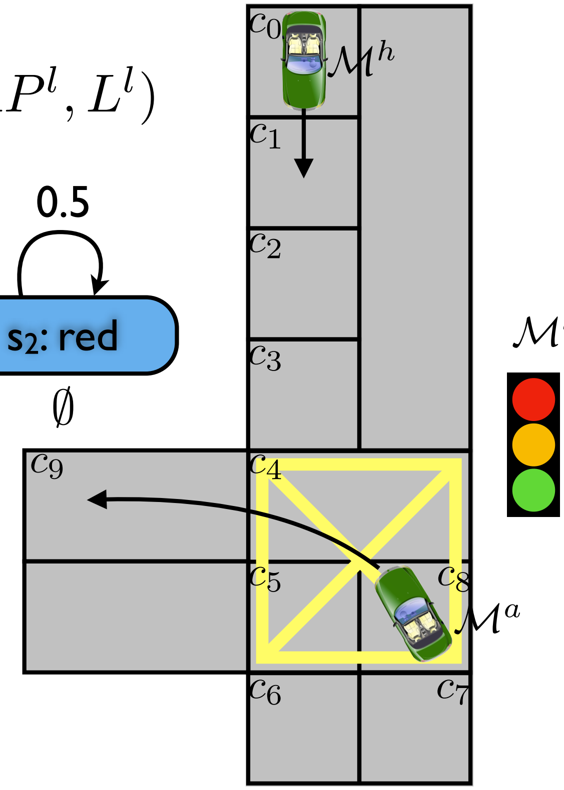

$\mathcal{M}^a || \mathcal{M}^l = (S^a \times S^l, Act, \mathbf{P}, \iota_{init}, AP^a \cup AP^l, L)$

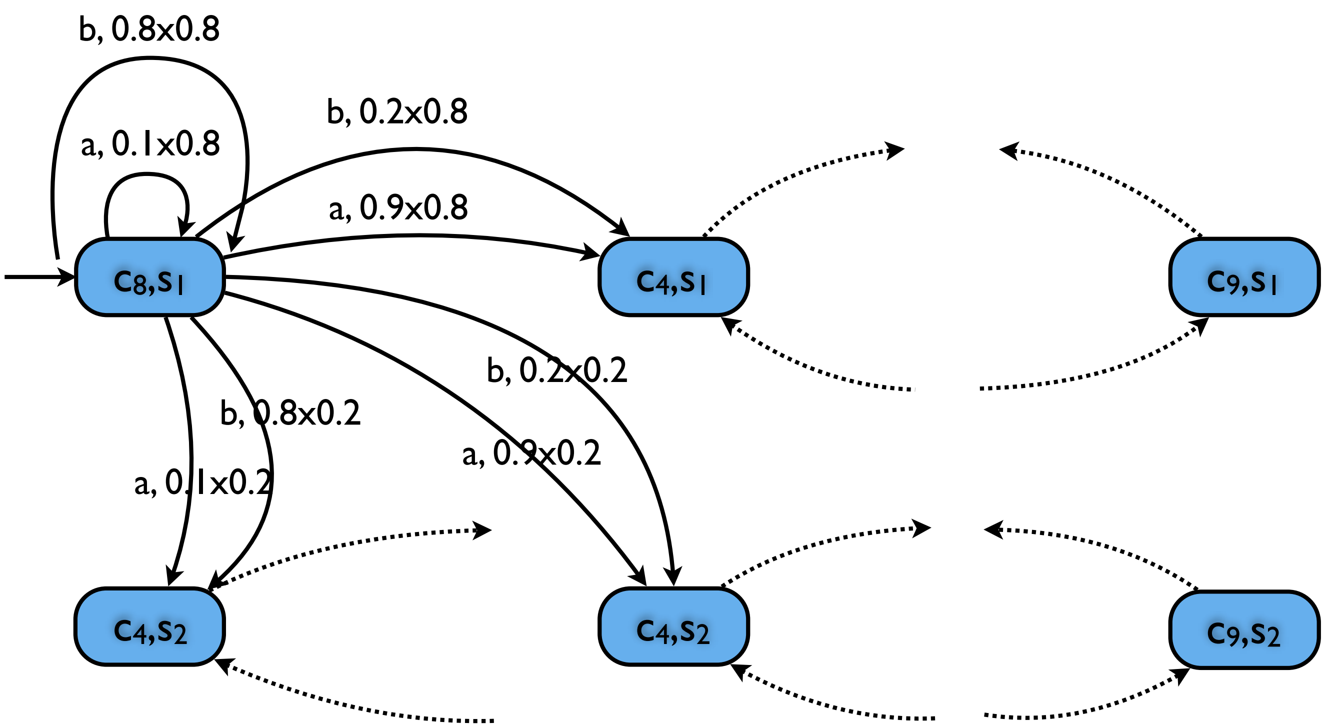

$\mathbf{P}\big((c, s), \alpha, (c', s')\big)$
$= \mathbf{P}^a(c, \alpha, c') \times \mathbf{P}^l(s, s')$

$\iota_{init}(c, s)$
$= \iota^a_{init}(c) \times \iota^l_{init}(s)$

$L(c, s) = L^a(c) \cup L^l(s)$

A policy of an MDP

Consider an MDP $\mathcal{M} = (S, Act, \mathbf{P}, \iota_{init}, AP, L)$.

- A *policy* for $\mathcal{M}$ is a function $\mathcal{C} : S^+ \to Act$ such that $\mathcal{C}(s_0 s_1 \dots s_n) \in Act(s_n)$ for all $s_0 s_1 \dots s_n \in S^+$.
- A $\mathcal{C}$*-path fragment* is an infinite sequence $\pi = s_0 s_1 s_2 \dots$ on $\mathcal{M}$ generated under policy $\mathcal{C}$ if $\mathbf{P}(s_i, \mathcal{C}(s_0 s_1 \dots s_i), s_{i+1}) > 0$ for all i.
- A policy $\mathcal{C}$ resolves all the nondeterministic choices in $\mathcal{M}$ and induces a Markov chain $\mathcal{M}_{\mathcal{C}} = \left(S^+, \mathbf{P}_{\mathcal{C}}, \iota_{init}, AP, L'\right)$ where for $\sigma = s_0 s_1 \dots s_n$,

$$\begin{aligned} \mathbf{P}_{\mathcal{C}}(\sigma, \sigma s_{n+1}) &= \mathbf{P}(s_n, \mathcal{C}(\sigma), s_{n+1}) \\ L'(\sigma) &= L(s_n) \end{aligned}$$

- $\mathcal{M}_{\mathcal{C}}$ is infinite even if $\mathcal{M}$ is finite.

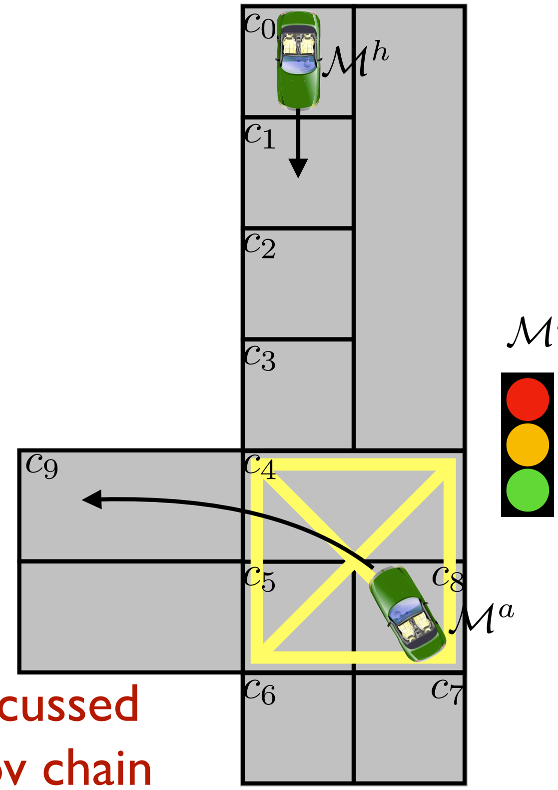

All we've discussed about Markov chain applies to $\mathcal{M}_{\mathcal{C}}$

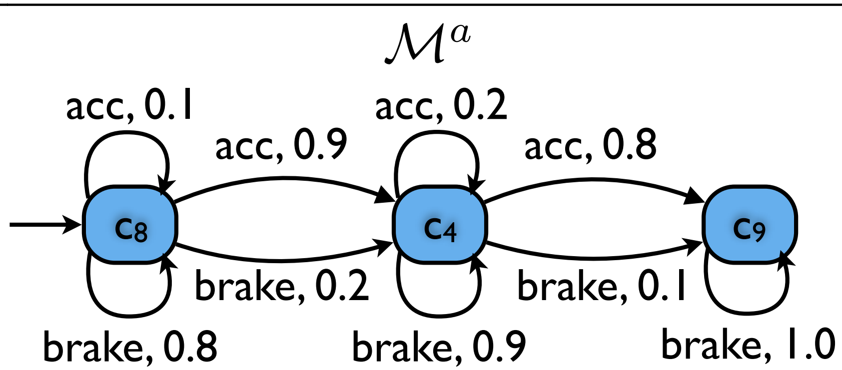

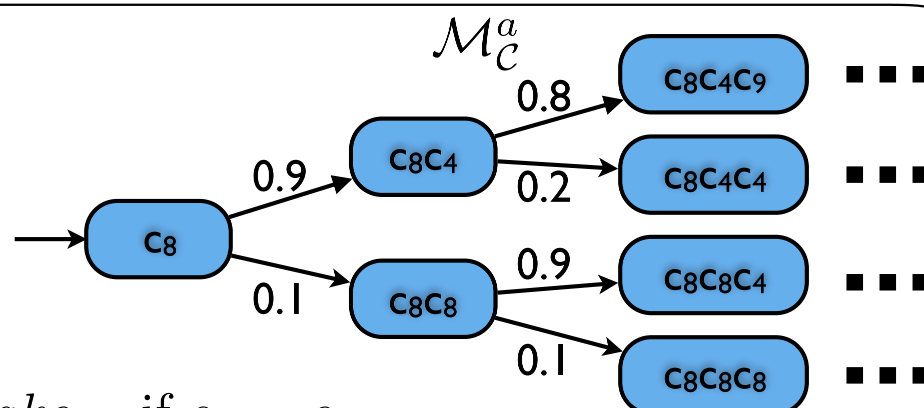

$$\mathcal{C}(s_0 \dots s_n) = \begin{cases} brake & \text{if } s_n = c_9 \\ acc & \text{otherwise} \end{cases}$$

Policy synthesis for MDPs with LTL specifications

Given an MDP $\mathcal{M} = (S, Act, \mathbf{P}, \iota_{init}, AP, L)$ and an LTL specification φ, compute an optimal policy $\mathcal{C}$ such that

$$Pr^{\mathcal{C}}(\varphi) = Pr^{\mathcal{M}_{\mathcal{C}}}(\varphi) = \sup_{\mathcal{C}'} Pr^{\mathcal{C}'}(\varphi)$$

An *end component* of $\mathcal{M}$ is a pair (T, A) where $\emptyset \neq T \subseteq S$ and $A : T \to 2^{Act}$ such that

- $0 \neq A(s) \subseteq Act(s)$ for all $s \in T$
- the directed graph induced by (T, A) is strongly connected
- for all $s \in T$ and $\alpha \in A(s)$,

$$\{t \in S \mid \mathbf{P}(s, \alpha, t) > 0\} \subseteq T.$$

Starting from any state in T, there exists a finite memory policy for $\mathscr{M}$ to keep the state within T forever while visiting all states in T infinitely often with probability 1.

$\Rightarrow$ At each state $s \in T$, select an action $\alpha \in A(s)$ according to a round-robin policy.

An end component is *maximal* if there is no end component $(T', A') \neq (T, A)$ such that $T \subseteq T'$ and $A(s) \subseteq A'(s)$ for all $s \in T$.

Reachability property

Consider an MDP $\mathcal{M} = (S, Act, \mathbf{P}, \iota_{init}, AP, L)$ and $B \subseteq S$.

For each $s \in S$, define $x_s = Pr_{\max}(s \models \Diamond B) = \sup_{\mathcal{C}} Pr^{\mathcal{C}}(s \models \Diamond B)$.

- If B is not reachable from s, $x_s = 0$.
- $x_s = 1$ for all $s \in B$.
- Define $\tilde{S} = \{s \in S \setminus B \mid B \text{ is reachable from } s\}$. For any $s \in \tilde{S}$,

$$\boxed{x_s \quad = \quad \max\left\{\sum_{t \in S} \mathbf{P}(s, \alpha, t) \cdot x_t \mid \alpha \in Act(s)\right\}}$$

Value iteration: $x_s^{(0)} = 0$ and $x_s^{(n+1)} = \max\left\{\sum_{t \in S} \mathbf{P}(s, \alpha, t) \cdot x_t^{(n)} \mid \alpha \in Act(s)\right\}$

Linear program: $\min \sum_{s \in S} x_s$ such that $x_s \geq \sum_{t \in S} \mathbf{P}(s, \alpha, t) \cdot x_t$ for all $\alpha \in Act$

DRA-based policy synthesis

MDP DRA

$$\mathcal{M} = (S, Act, \mathbf{P}, \iota_{init}, AP, L) \quad \otimes \quad \mathcal{A} = (Q, 2^{AP}, \delta, q_0, Acc)$$

$$= (S', Act, \mathbf{P}', \iota'_{init}, AP', L')$$

Product MDP

- S', ι'_{init}, AP', L' are defined as in the product of Markov chain and DRA.
- $\mathbf{P}'(\langle s, q\rangle, \alpha, \langle s', q'\rangle) = \begin{cases} \mathbf{P}(s, \alpha, s') & \text{if } q' = \delta(q, L(s')) \\ 0 & \text{otherwise} \end{cases}$
- For each $(L_i, K_i) \in Acc$, let $\mathcal{M}_{\Box\neg L_i}$ be the MDP that results from $\mathcal{M} \otimes \mathcal{A}$ by removing all the states in $S \times L_i$ and removing all the actions $\alpha \in Act(s)$ such that $Post(s, \alpha) \subseteq S \times L_i$. Define

$$\boxed{U_i = \bigcup_{\substack{(T, A) \in MEC(\mathcal{M}_{\Box\neg L_i}) \\ T \cap (S \times K_i) \neq \emptyset}} T}$$

$U = \cup_i U_i$ is known as the *success set* for Rabin Conditions

$$\boxed{Pr^{\mathcal{M}}_{max}(s \models \varphi) = Pr^{\mathcal{M}\otimes\mathcal{A}}_{max}\big(\langle s, \delta(q_0, L(s))\rangle \models \Diamond U\big)}$$