

# Chapter 11


## Acting with Nondeterministic Models

Dana S. Nau  
University of Maryland



# Motivation

- In Chapters 8–10, we had probability distributions over the possible outcomes of actions
- Sometimes we want to reason about nondeterminism *without* the probability distributions
  - ▶ Probabilities might not be available
  - ▶ We might want policies that satisfy safety conditions:
    - Guaranteed to work for all possible action outcomes



Credit: [Dennis Hill](#), [CC BY 2.0](#)



Credit: [Airtuna08](#), [CC BY-SA 3.0](#)




Credit: [David Wilson](#), [CC BY 2.0](#)



Credit: [Slaunger, CC BY-SA 3.0](#)

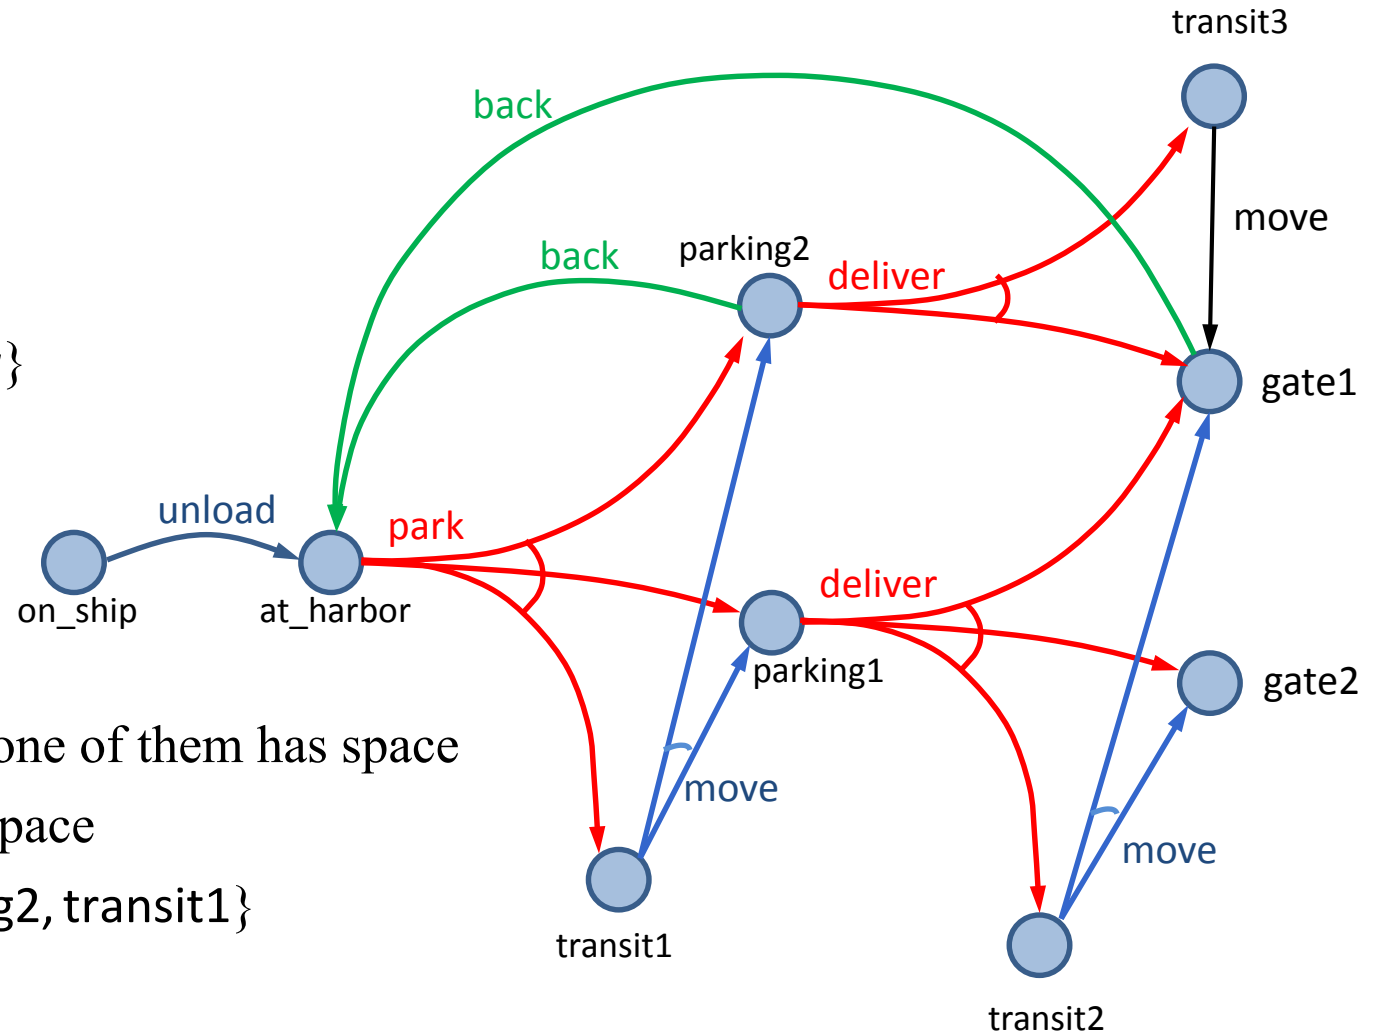
# Example

- Very simple harbor management domain
  - ▶ Unload a single item from a ship
  - ▶ Move it around a harbor
- One state variable:  $\text{pos}(\text{item})$ 
  - ▶ Simplified names for states
  - ▶ For  $\{\text{pos}(\text{item})=\text{on\_ship}\}$ , just write  $\text{on\_ship}$



# Nondeterministic Planning Domains

- 3-tuple  $(S, A, \gamma)$ 
  - ▶  $S$  and  $A$  – finite sets of states and actions
  - ▶  $\gamma: S \times A \rightarrow 2^S$
- $\gamma(s, a) = \{\text{all possible “next states” after applying action } a \text{ in state } s\}$ 
  - ▶  $a$  is applicable in state  $s$  iff  $\gamma(s, a) \neq \emptyset$
- $\text{Applicable}(s) = \{\text{all actions applicable in } s\}$   
 $= \{a \in A \mid \gamma(s, a) \neq \emptyset\}$
- Example:
  - ▶  $\text{Applicable}(\text{at\_harbor}) = \{\text{park}\}$
  - ▶ park has three possible outcomes
    - put item in parking1 or parking2 if one of them has space
    - or in transit1 if there’s no parking space
  - ▶  $\gamma(\text{at\_harbor}, \text{park}) = \{\text{parking1}, \text{parking2}, \text{transit1}\}$



# Nondeterministic Planning Domains

- One possible action representation:
  - ▶ like classical, but with  $n$  mutually exclusive “effects” lists

- e.g., park:

pre:  $\text{pos}(\text{item}) = \text{at\_harbor}$

eff<sub>1</sub>:  $\text{pos}(\text{item}) \leftarrow \text{parking1}$

eff<sub>2</sub>:  $\text{pos}(\text{item}) \leftarrow \text{parking2}$

eff<sub>3</sub>:  $\text{pos}(\text{item}) \leftarrow \text{transit1}$

- Problem:

- ▶ number of effects lists may be combinatorially large

- ▶ Suppose  $a$  can cause any possible combination of effects  $e_1, e_2, \dots, e_k$

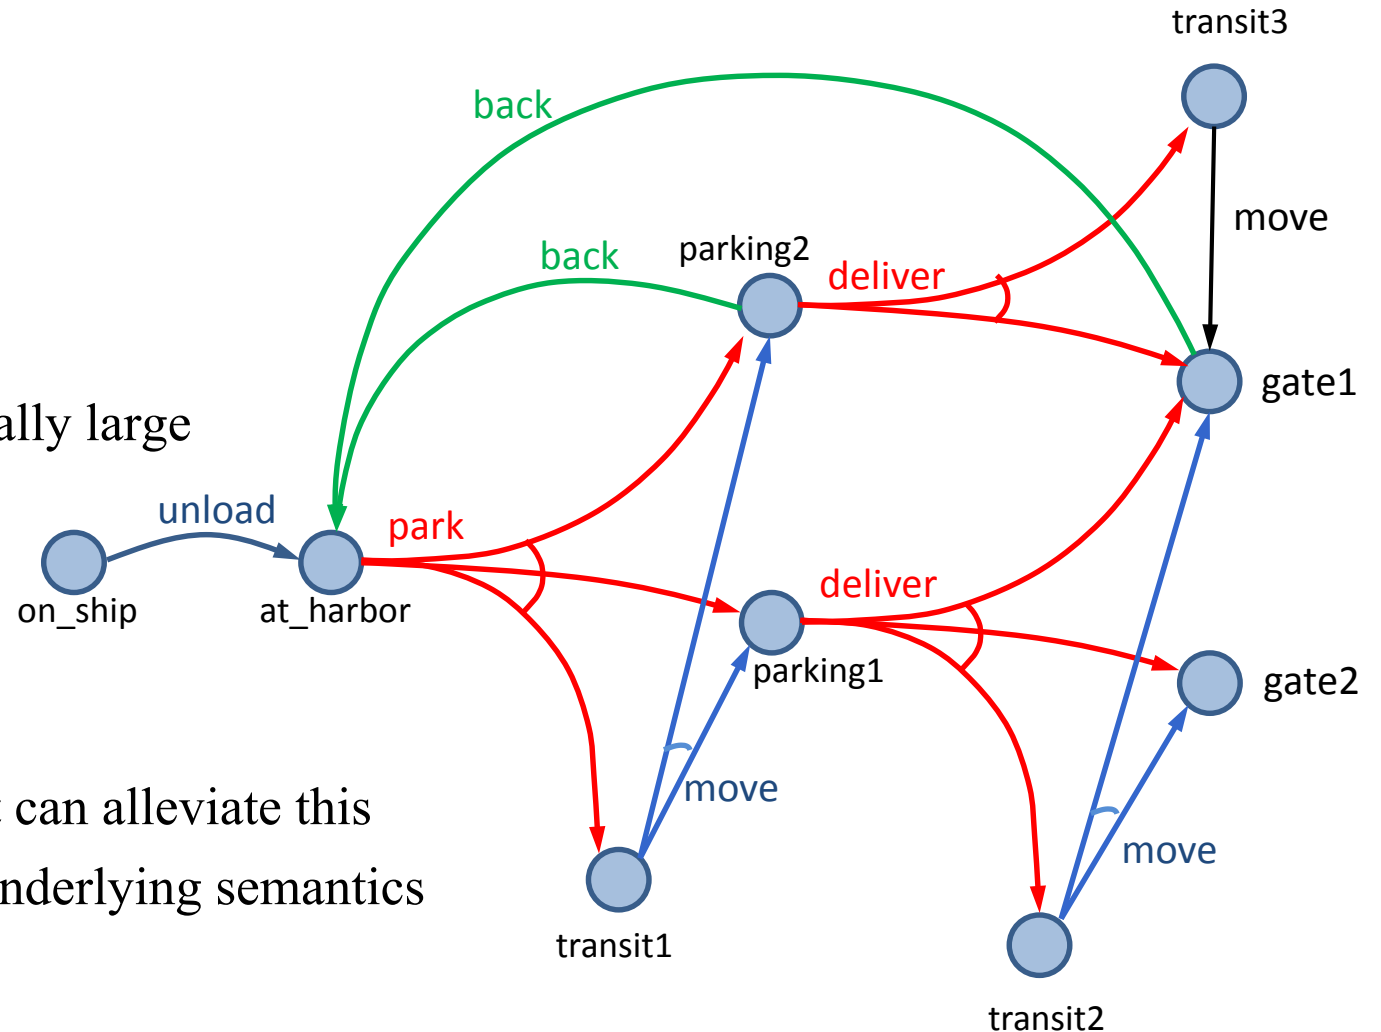
- ▶ Need  $\text{eff}_1, \text{eff}_2, \dots, \text{eff}_{2^k}$

- One for each combination

- ▶ Section 12.3: a different representation that can alleviate this


- For now, ignore most of that, just look at the underlying semantics

- ▶ states, actions  $\Leftrightarrow$  nodes, edges in a graph



# Nondeterministic Planning Domains

- For deterministic planning problems, search space was a graph
- Now it's an AND/OR graph
  - ▶ *OR branch*:
    - several applicable actions, which one to choose?
  - ▶ *AND branch*:
    - multiple possible outcomes, must handle all of them
- Analogy to PSP in Chapter 2
  - ▶ *OR branch*  $\Leftrightarrow$  action selection
  - ▶ *AND branch*  $\Leftrightarrow$  flaw selection



# Policies

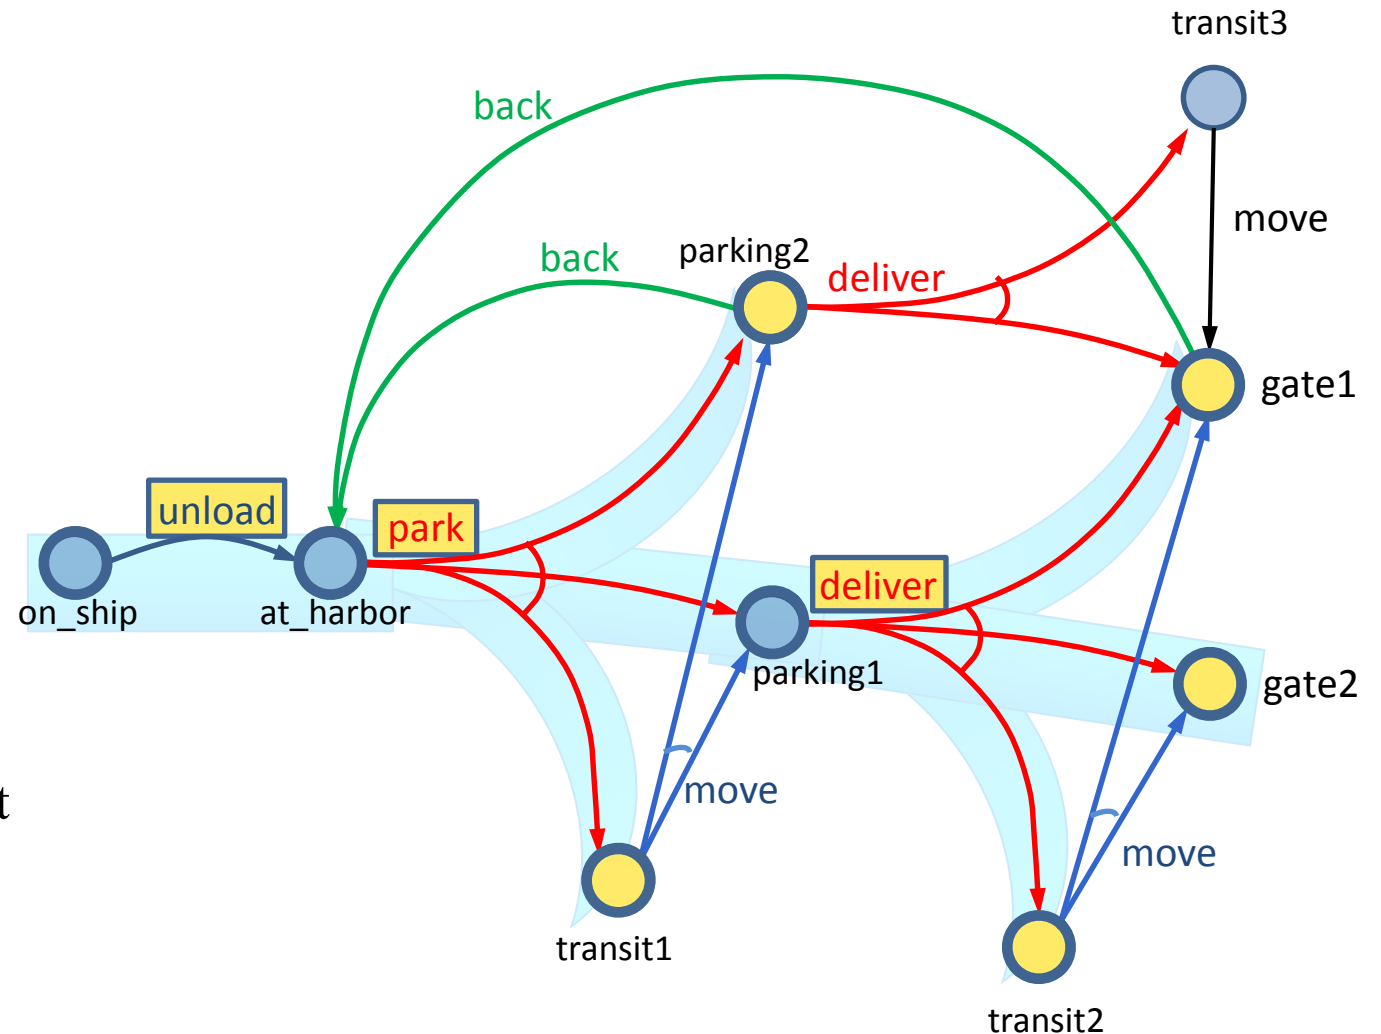
- *Policy*: a function  $\pi : S' \rightarrow A$ 
  - ▶  $S' \subseteq S$
  - ▶ For every  $s \in \text{Domain}(\pi)$ , require  $\pi(s) \in \text{Applicable}(s)$

- Two equivalent notations:

$\pi_1(\text{on\_ship}) = \text{unload}$ ,  
 $\pi_1(\text{at\_harbor}) = \text{park}$ ,  
 $\pi_1(\text{parking1}) = \text{deliver}$

$\pi_1 = \{(\text{on\_ship}, \text{unload}),$   
 $(\text{at\_harbor}, \text{park}),$   
 $(\text{parking1}, \text{deliver})\}$

- That's just the notation
  - ▶ implementation could be quite different



# Definitions Over Policies

- *Transitive closure:*

- ▶  $\hat{\gamma}(s, \pi) = \{\text{all states reachable from } s \text{ using } \pi\}$

- ▶  $\hat{\gamma}(s, \pi) = S_0 \cup S_1 \cup S_2 \cup \dots$

$$S_0 = \{s\}$$

$$S_1 = S_0 \cup \{\gamma(s_0, \pi(s_0)) \mid s_0 \in S_0\}$$

$$S_2 = S_1 \cup \{\gamma(s_1, \pi(s_1)) \mid s_1 \in S_1\}$$

...

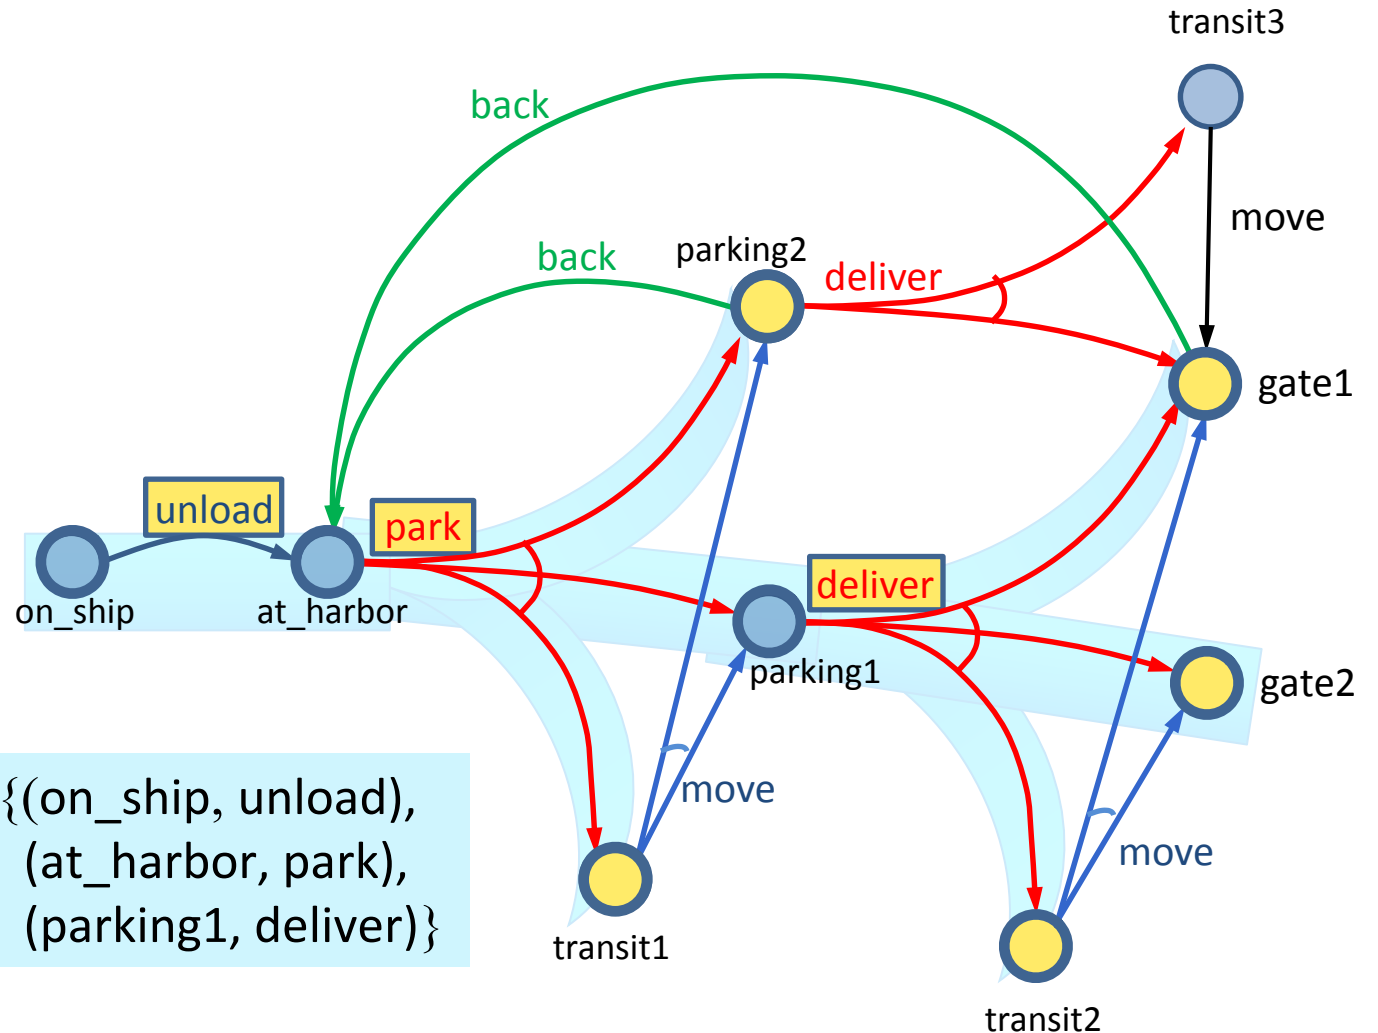
- *Reachability graph:*  $\text{Graph}(s, \pi) = (V, E)$

- ▶  $V = \hat{\gamma}(s, \pi)$

- ▶  $E = \{(s_1, s_2) \mid s_1 \in V, s_2 \in \gamma(s_1, \pi(s_1))\}$

- $\text{leaves}(s, \pi) = \hat{\gamma}(s, \pi) \setminus \text{Dom}(\pi)$

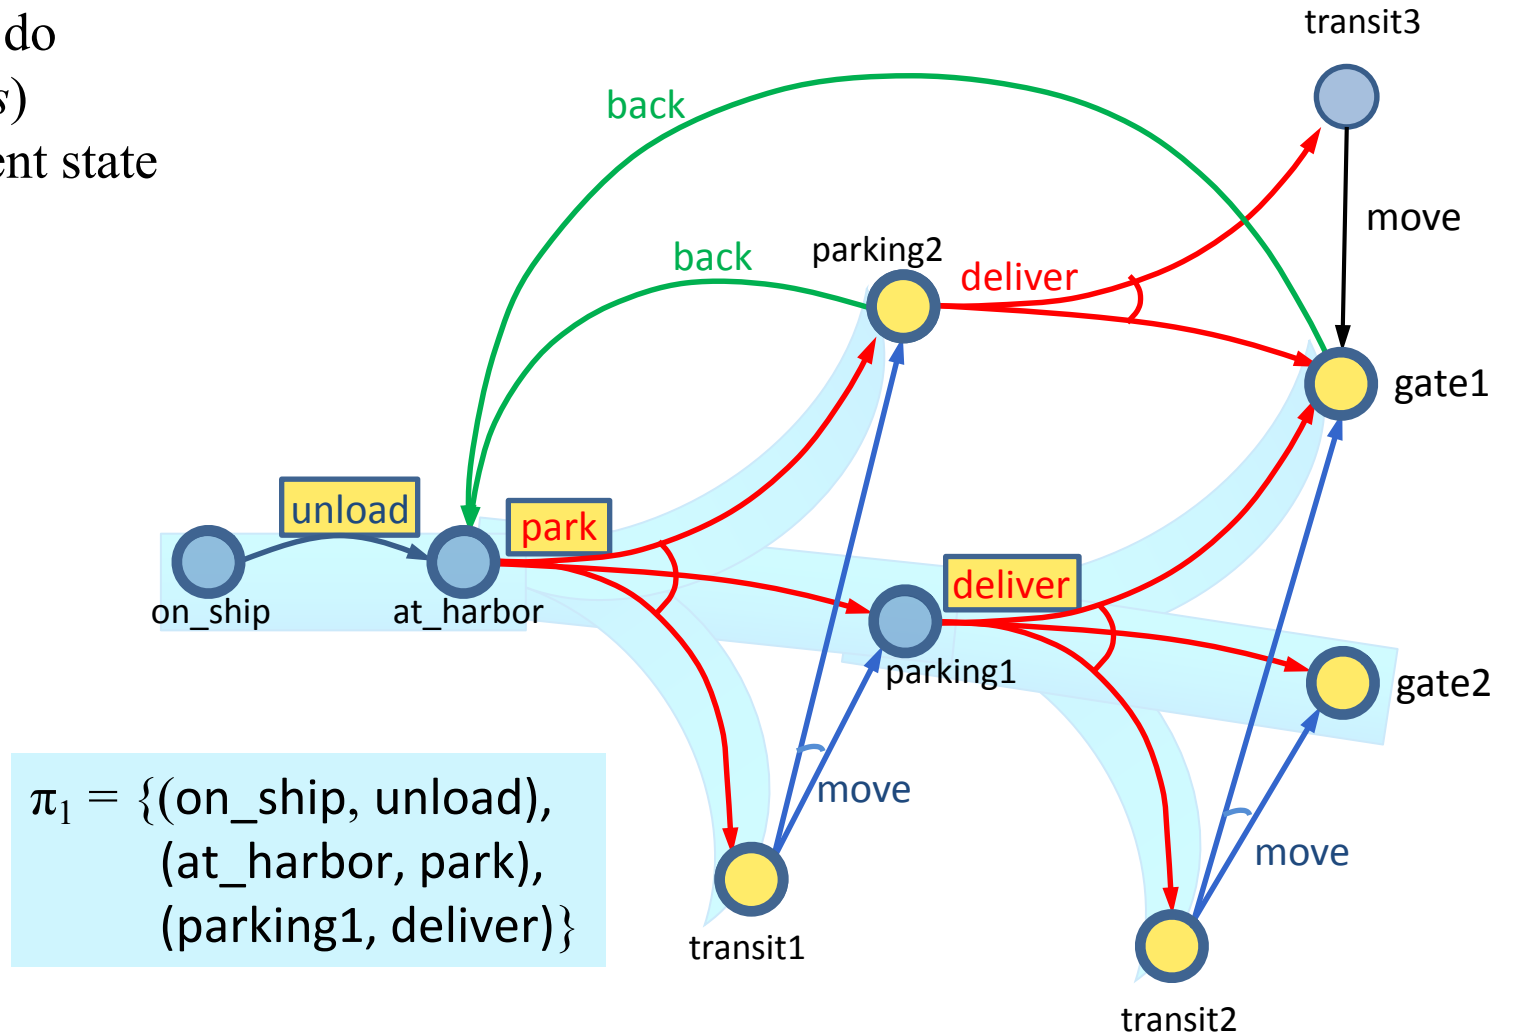
- ▶ may be empty





# Acting with a Policy

- ActPolicy( $\pi$ )
  - $s \leftarrow$  observe current state
  - while  $s \in \text{Domain}(\pi)$  do
  - perform action  $\pi(s)$
  - $s \leftarrow$  observe current state

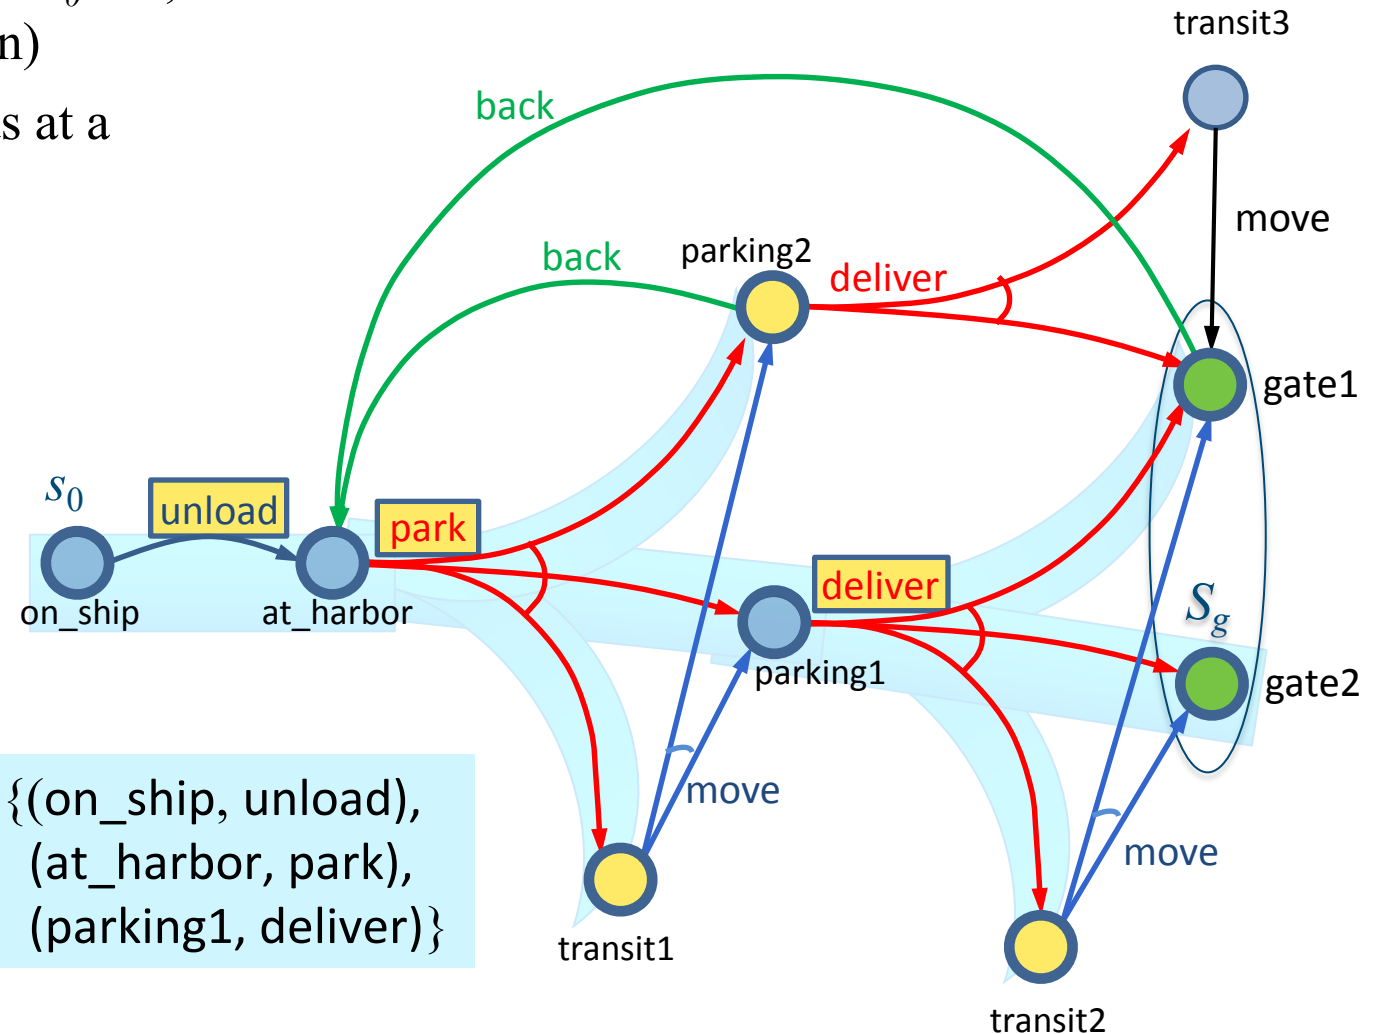


# Types of Policies

- Acting (or planning) problem  $P = (\Sigma, s_0, S_g)$ 
  - planning domain  $\Sigma = (S, A, \gamma)$ , initial state  $s_0 \in S$ , set of goal states  $S_g \subseteq S$  (shown in green)
- $\pi$  is a *solution* if at least one execution ends at a goal
  - $leaves(s, \pi) \cap S_g \neq \emptyset$
- A policy  $\pi$  is *safe* if
  - $\forall s \in \hat{\gamma}(s_0, \pi), leaves(s, \pi) \cap S_g \neq \emptyset$
  - at every state in  $\hat{\gamma}(s_0, \pi)$ , at least one of the execution paths from  $s$  using  $\pi$  stops at a goal state.
- Otherwise, *unsafe* policy

**Poll:** Is  $\pi_1$  safe or unsafe?

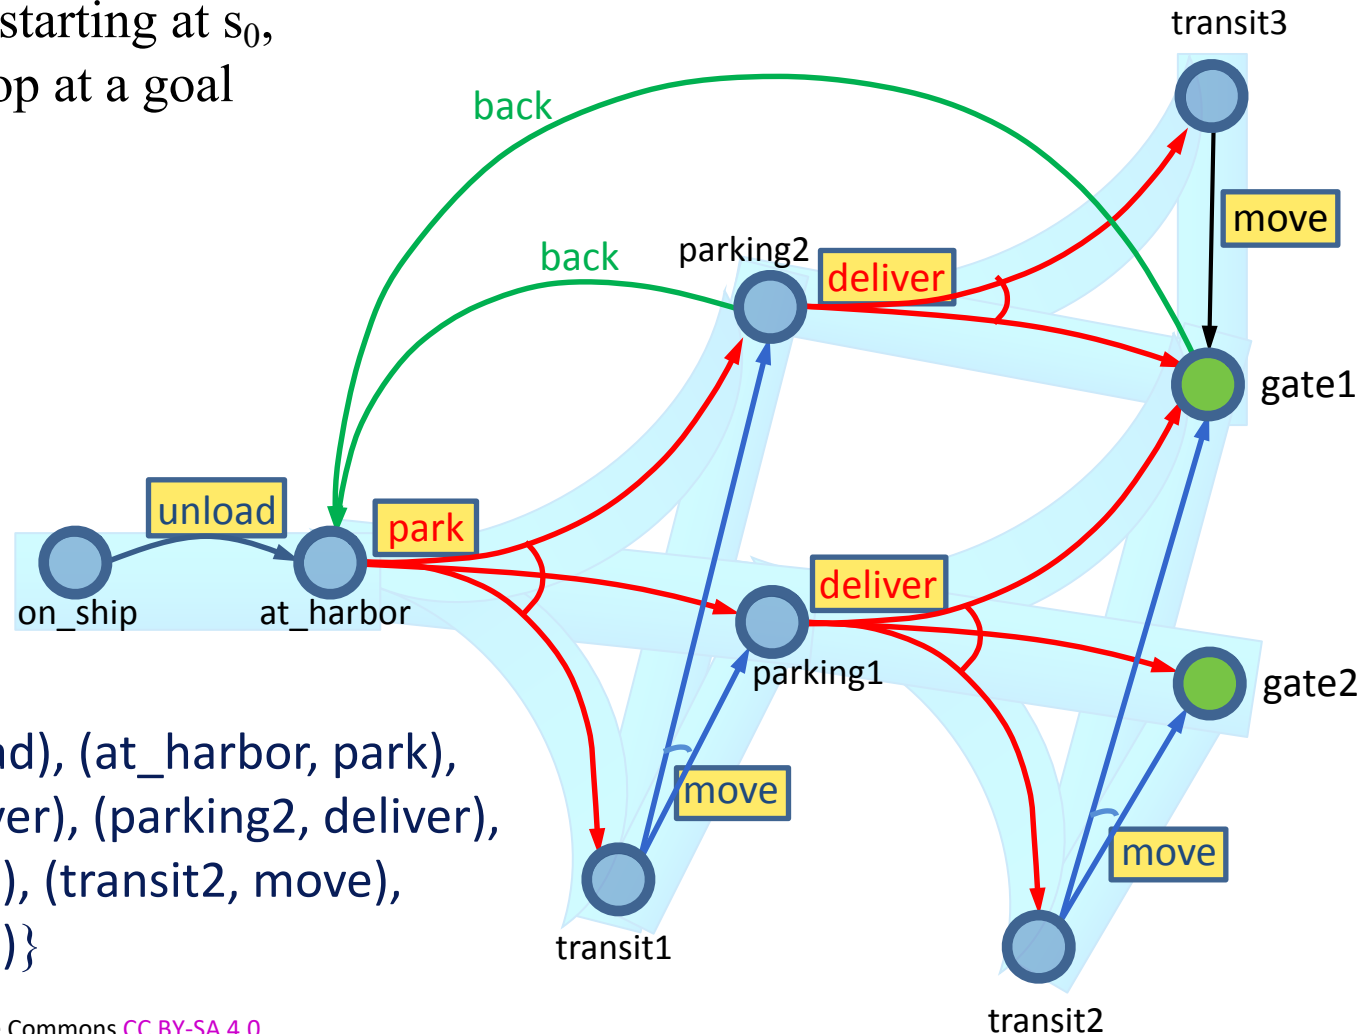
$\pi_1 = \{(on\_ship, unload), (at\_harbor, park), (parking1, deliver)\}$



# Safe Policies

- *Acyclic* safe policy
  - $\text{Graph}(s_0, \pi)$  is acyclic, and  $\text{leaves}(s, \pi) \subseteq S_g$
- If we run  $\text{ActPolicy}(\pi)$  starting at  $s_0$ , we're guaranteed to stop at a goal

- $\text{ActPolicy}(\pi)$ 
  - $s \leftarrow$  observe current state
  - while  $s \in \text{Domain}(\pi)$  do
  - perform action  $\pi(s)$
  - $s \leftarrow$  observe current state



$\pi_2 = \{(\text{on\_ship}, \text{unload}), (\text{at\_harbor}, \text{park}),$   
 $(\text{parking1}, \text{deliver}), (\text{parking2}, \text{deliver}),$   
 $(\text{transit1}, \text{move}), (\text{transit2}, \text{move}),$   
 $(\text{transit3}, \text{move})\}$


# Safe Policies

- *Cyclic* safe policy

- ▶  $\text{Graph}(s_0, \pi)$  is cyclic, and  $\text{leaves}(s, \pi) \subseteq S_g$ , and  $\forall s \in \hat{\gamma}(s_0, \pi), \text{leaves}(s, \pi) \cap S_g \neq \emptyset$

- At every state  $s$  in  $\hat{\gamma}(s_0, \pi)$ , at least one of the execution paths from  $s$  using  $\pi$  ends at a goal state
- ▶ Will never get caught in a dead end

$\pi_3 = \{(\text{on\_ship}, \text{unload}), (\text{at\_harbor}, \text{park}), (\text{parking1}, \text{deliver}), (\text{parking2}, \text{back}), (\text{transit1}, \text{move}), (\text{transit2}, \text{move}), (\text{gate1}, \text{back})\}$



- $\text{ActPolicy}(\pi)$ 
  - $s \leftarrow$  observe current state
  - while  $s \in \text{Domain}(\pi)$  do
  - perform action  $\pi(s)$
  - $s \leftarrow$  observe current state

**Poll:** Let  $\pi$  be a cyclic safe solution. Suppose we run  $\text{ActPolicy}(\pi)$  starting at  $s_0$ .

1. Are there situations where we can be sure  $\pi$  will reach a goal?
2. Are there situations where we can't be sure  $\pi$  will reach a goal?

# Safe Policies

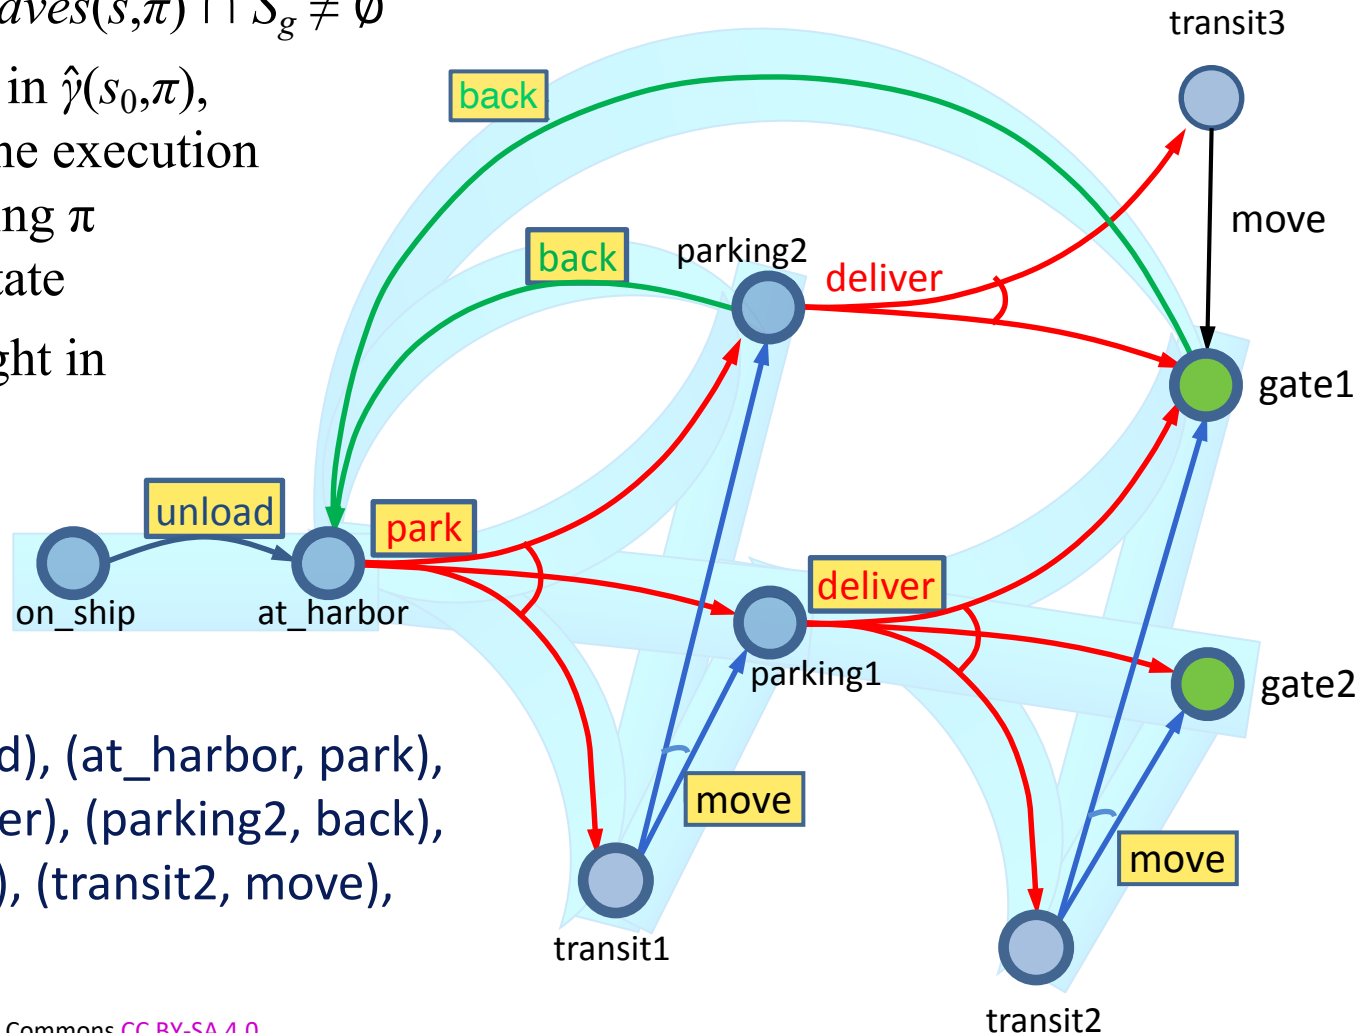
- *Cyclic* safe policy

- ▶  $\text{Graph}(s_0, \pi)$  is cyclic, and  $\text{leaves}(s, \pi) \subseteq S_g$ , and  $\forall s \in \hat{\gamma}(s_0, \pi), \text{leaves}(s, \pi) \cap S_g \neq \emptyset$

- At every state  $s$  in  $\hat{\gamma}(s_0, \pi)$ , at least one of the execution paths from  $s$  using  $\pi$  ends at a goal state

- ▶ Will never get caught in a dead end
- ▶ Every “fair” execution will reach a goal

$\pi_3 = \{(\text{on\_ship}, \text{unload}), (\text{at\_harbor}, \text{park}), (\text{parking1}, \text{deliver}), (\text{parking2}, \text{back}), (\text{transit1}, \text{move}), (\text{transit2}, \text{move}), (\text{gate1}, \text{back})\}$




- $\text{ActPolicy}(\pi)$

$s \leftarrow$  observe current state  
 while  $s \in \text{Domain}(\pi)$  do  
     perform action  $\pi(s)$   
 $s \leftarrow$  observe current state

**Poll:**

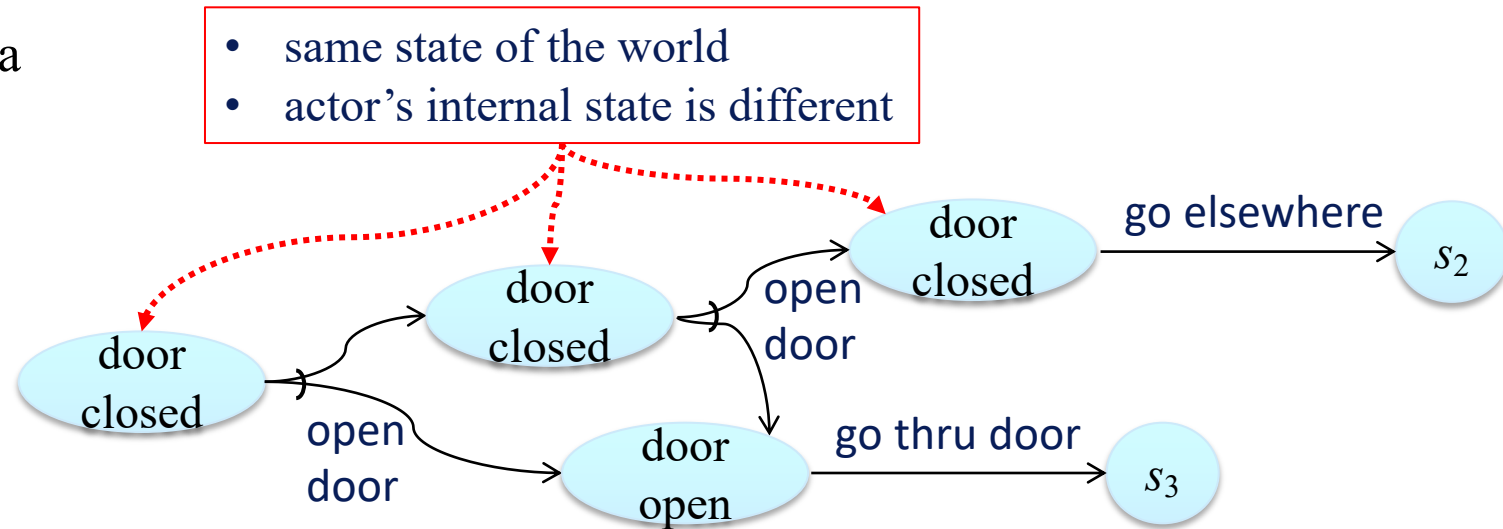
1. Can you think of a real-world situation in which all executions are “fair”?
2. Can you think of a real-world situation in which there are “unfair” executions?

# Kinds of Solution Policies



# Beyond Policies

- Sometimes we want to give the actor instructions that can't be described as a policy, e.g.,
  - Try to open the door twice.  
If it opens, go through it.  
If it doesn't, go to another door
- The book describes two other ways to represent instructions to the actor
  - ▶ Input/Output Automata
  - ▶ Behavior Trees
- I'll discuss behavior trees in a separate set of slides



# Summary

- Actions, plans, policies, planning problems
- Types of solution policies:
  - ▶ unsafe, safe (acyclic, cyclic)
- Motivation for instructions other than policies