

State-Space Search and the STRIPS Planner

Searching for a Path
through a Graph of Nodes
Representing World States

Literature

- Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated Planning – Theory and Practice*, chapter 2 and 4. Elsevier/Morgan Kaufmann, 2004.
- Malik Ghallab, *et al.* PDDL–The Planning Domain Definition Language, Version 1.x.
<ftp://ftp.cs.yale.edu/pub/mcdermott/software/pddl.tar.gz>
- S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*, chapters 3-4. Prentice Hall, 2nd edition, 2003.
- J. Pearl. *Heuristics*, chapters 1-2. Addison-Wesley, 1984.

Classical Representations

- propositional representation
 - world state is set of propositions
 - action consists of precondition propositions, propositions to be added and removed
- STRIPS representation
 - like propositional representation, but first-order literals instead of propositions
- state-variable representation
 - state is tuple of state variables $\{x_1, \dots, x_n\}$
 - action is partial function over states

Overview

- **The STRIPS Representation**
- The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- Heuristic Search
- Forward State-Space Search
- Backward State-Space Search
- The STRIPS Planner

STRIPS Planning Domains: Restricted State-Transition Systems


- A restricted state-transition system is a triple $\Sigma=(S,A,\gamma)$, where:
 - $S=\{s_1,s_2,\dots\}$ is a set of states;
 - $A=\{a_1,a_2,\dots\}$ is a set of actions;
 - $\gamma:S\times A\rightarrow S$ is a state transition function.
- defining STRIPS planning domains:
 - define STRIPS states
 - define STRIPS actions
 - define the state transition function

States in the STRIPS Representation

- Let \mathcal{L} be a first-order language with finitely many predicate symbols, finitely many constant symbols, and no function symbols.
- A state in a STRIPS planning domain is a set of ground atoms of \mathcal{L} .
 - (ground) atom p holds in state s iff $p\in s$
 - s satisfies a set of (ground) literals g (denoted $s \models g$) if:
 - every positive literal in g is in s and
 - every negative literal in g is not in s .

DWR Example: STRIPS States

```
state = {attached(p1,loc1),
        attached(p2,loc1),
        in(c1,p1),in(c3,p1),
        top(c3,p1), on(c3,c1),
        on(c1,pallet), in(c2,p2),
        top(c2,p2), on(c2,pallet),
        belong(crane1,loc1),
        empty(crane1),
        adjacent(loc1,loc2),
        adjacent(loc2, loc1),
        at(r1,loc2), occupied(loc2),
        unloaded(r1)}
```



Fluent Relations

- Predicates that represent relations, the truth value of which can change from state to state, are called a fluent or flexible relations.
 - example: `at`
- A state-invariant predicate is called a rigid relation.
 - example: `adjacent`

Operators and Actions in STRIPS Planning Domains

- A planning operator in a STRIPS planning domain is a triple $o = (\text{name}(o), \text{precond}(o), \text{effects}(o))$ where:
 - the name of the operator $\text{name}(o)$ is a syntactic expression of the form $n(x_1, \dots, x_k)$ where n is a (unique) symbol and x_1, \dots, x_k are all the variables that appear in o , and
 - the preconditions $\text{precond}(o)$ and the effects $\text{effects}(o)$ of the operator are sets of literals.
- An action in a STRIPS planning domain is a ground instance of a planning operator.

DWR Example: STRIPS Operators

- $\text{move}(r, l, m)$
 - precondition: $\text{adjacent}(l, m), \text{at}(r, l), \neg \text{occupied}(m)$
 - effects: $\text{at}(r, m), \text{occupied}(m), \neg \text{occupied}(l), \neg \text{at}(r, l)$
- $\text{load}(k, l, c, r)$
 - precondition: $\text{belong}(k, l), \text{holding}(k, c), \text{at}(r, l), \text{unloaded}(r)$
 - effects: $\text{empty}(k), \neg \text{holding}(k, c), \text{loaded}(r, c), \neg \text{unloaded}(r)$
- $\text{put}(k, l, c, d, p)$
 - precondition: $\text{belong}(k, l), \text{attached}(p, l), \text{holding}(k, c), \text{top}(d, p)$
 - effects: $\neg \text{holding}(k, c), \text{empty}(k), \text{in}(c, p), \text{top}(c, p), \text{on}(c, d), \neg \text{top}(d, p)$

Applicability and State Transitions

- Let L be a set of literals.
 - L^+ is the set of atoms that are positive literals in L and
 - L^- is the set of all atoms whose negations are in L .
- Let a be an action and s a state. Then a is applicable in s iff:
 - $\text{precond}^+(a) \subseteq s$; and
 - $\text{precond}^-(a) \cap s = \{\}$.
- The state transition function γ for an applicable action a in state s is defined as:
 - $\gamma(s,a) = (s - \text{effects}^-(a)) \cup \text{effects}^+(a)$

STRIPS Planning Domains



- Let \mathcal{L} be a function-free first-order language. A STRIPS planning domain on \mathcal{L} is a restricted state-transition system $\Sigma=(S,A,\gamma)$ such that:
 - S is a set of STRIPS states, i.e. sets of ground atoms
 - A is a set of ground instances of some STRIPS planning operators O
 - $\gamma: S \times A \rightarrow S$ where
 - $\gamma(s,a) = (s - \text{effects}^-(a)) \cup \text{effects}^+(a)$ if a is applicable in s
 - $\gamma(s,a) = \text{undefined}$ otherwise
 - S is closed under γ

STRIPS Planning Problems

- A STRIPS planning problem is a triple $\mathcal{P}=(\Sigma,s_i,g)$ where:
 - $\Sigma=(S,A,\gamma)$ is a STRIPS planning domain on some first-order language \mathcal{L}
 - $s_i \in S$ is the initial state
 - g is a set of ground literals describing the goal such that the set of goal states is: $S_g=\{s \in S \mid s \text{ satisfies } g\}$

DWR Example: STRIPS Planning Problem

- Σ : STRIPS planning domain for DWR domain
- s_i : any state
 - example: $s_0 = \{ \text{attached}(\text{pile}, \text{loc1}), \text{in}(\text{cont}, \text{pile}), \text{top}(\text{cont}, \text{pile}), \text{on}(\text{cont}, \text{pallet}), \text{belong}(\text{crane}, \text{loc1}), \text{empty}(\text{crane}), \text{adjacent}(\text{loc1}, \text{loc2}), \text{adjacent}(\text{loc2}, \text{loc1}), \text{at}(\text{robot}, \text{loc2}), \text{occupied}(\text{loc2}), \text{unloaded}(\text{robot}) \}$
- g : any subset of L
 - example: $g = \{ \text{-unloaded}(\text{robot}), \text{at}(\text{robot}, \text{loc2}) \}$, i.e. $S_g = \{s_5\}$



Statement of a STRIPS Planning Problem

- A statement of a STRIPS planning problem is a triple $P=(O, s_i, g)$ where:
 - O is a set of planning operators in an appropriate STRIPS planning domain $\Sigma=(S, A, \gamma)$ on \mathcal{L}
 - s_i is the initial state in an appropriate STRIPS planning problem $\mathcal{P}=(\Sigma, s_i, g)$
 - g is a goal (set of ground literals) in the same STRIPS planning problem \mathcal{P}

Classical Plans

- A plan is any sequence of actions $\pi=\langle a_1, \dots, a_k \rangle$, where $k \geq 0$.
 - The length of plan π is $|\pi|=k$, the number of actions.
 - If $\pi_1=\langle a_1, \dots, a_k \rangle$ and $\pi_2=\langle a'_1, \dots, a'_j \rangle$ are plans, then their concatenation is the plan $\pi_1 \bullet \pi_2 = \langle a_1, \dots, a_k, a'_1, \dots, a'_j \rangle$.
 - The extended state transition function for plans is defined as follows:
 - $\gamma(s, \pi) = s$ if $k=0$ (π is empty)
 - $\gamma(s, \pi) = \gamma(\gamma(s, a_1), \langle a_2, \dots, a_k \rangle)$ if $k > 0$ and a_1 applicable in s
 - $\gamma(s, \pi) = \text{undefined}$ otherwise

Classical Solutions

- Let $\mathcal{P}=(\Sigma, s_i, g)$ be a planning problem. A plan π is a solution for \mathcal{P} if $\gamma(s_i, \pi)$ satisfies g .
 - A solution π is redundant if there is a proper subsequence of π is also a solution for \mathcal{P} .
 - π is minimal if no other solution for \mathcal{P} contains fewer actions than π .

DWR Example: Solution Plan

- plan $\pi_1 =$
 - $\langle \text{move}(\text{robot}, \text{loc2}, \text{loc1}),$
 - $\text{take}(\text{crane}, \text{loc1}, \text{cont}, \text{pallet}, \text{pile}),$
 - $\text{load}(\text{crane}, \text{loc1}, \text{cont}, \text{robot}),$
 - $\text{move}(\text{robot}, \text{loc1}, \text{loc2}) \rangle$
- $|\pi_1|=4$
- π_1 is a minimal, non-redundant solution

Overview

- The STRIPS Representation
- ◆ The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- Heuristic Search
- Forward State-Space Search
- Backward State-Space Search
- The STRIPS Planner

PDDL Basics

- <http://cs-www.cs.yale.edu/homes/dvm/>
- language features (version 1.x):
 - basic STRIPS-style actions
 - various extensions as explicit requirements
- used to define:
 - planning domains: requirements, types, predicates, possible actions
 - planning problems: objects, rigid and fluent relations, initial situation, goal description

PDDL 1.x Domains

```
<domain> ::=
  (define (domain <name>)
    [<extension-def>]
    [<require-def>]
    [<types-def>]:typing
    [<constants-def>]
    [<domain-vars-def>]:expression-evaluation
    [<predicates-def>]
    [<timeless-def>]
    [<safety-def>]:safety-constraints
    <structure-def>*)

<extension-def> ::=
  (.extends <domain name>+)
<require-def> ::=
  (.requirements <require-key>+)
<require-key> ::=
  :strips | :typing | ...

<types-def> ::= (:types <typed list (name)>)
<constants-def> ::=
  (:constants <typed list (name)>)
<domain-vars-def> ::=
  (:domain-variables
   <typed list (domain-var-declaration)>)
<predicates-def> ::=
  (:predicates <atomic formula skeleton>+)
<atomic formula skeleton> ::=
  (<predicate> <typed list (variable)>)
<predicate> ::= <name>
<variable> ::= ?<name>
<timeless-def> ::=
  (:timeless <literal (name)>+)
<structure-def> ::= <action-def>
<structure-def> ::= :domain-axioms <axiom-def>
<structure-def> ::= :action-expansions <method-def>
```

PDDL Types

- PDDL types syntax
 - <typed list (x)> ::= x*
 - <typed list (x)> ::=:typing
x+ - <type> <typed list(x)>
 - <type> ::= <name>
 - <type> ::= (either <type>+)
 - <type> ::=:fluents (fluent <type>)

Example: DWR Types

```
(define (domain dock-worker-robot)

  (:requirements :strips :typing )

  (:types
    location ;there are several connected locations
    pile ;is attached to a location,
           ;it holds a pallet and a stack of containers
    robot ;holds at most 1 container,
           ;only 1 robot per location
    crane ;belongs to a location to pickup containers
    container )

  ...)
```

Example: DWR Predicates

```
(:predicates
  (adjacent ?l1 ?l2 - location) ;location ?l1 is adjacent to ?l2
  (attached ?p - pile ?l - location) ;pile ?p attached to location ?l
  (belong ?k - crane ?l - location) ;crane ?k belongs to location ?l

  (at ?r - robot ?l - location) ;robot ?r is at location ?l
  (occupied ?l - location) ;there is a robot at location ?l
  (loaded ?r - robot ?c - container ) ;robot ?r is loaded with container ?c
  (unloaded ?r - robot) ;robot ?r is empty

  (holding ?k - crane ?c - container) ;crane ?k is holding a container ?c
  (empty ?k - crane) ;crane ?k is empty

  (in ?c - container ?p - pile) ;container ?c is within pile ?p
  (top ?c - container ?p - pile) ;container ?c is on top of pile ?p
  (on ?c1 - container ?c2 - container) ;container ?c1 is on container ?c2
)
```

PDDL Actions

```
<action-def> ::=
  (:action <action functor>
   :parameters ( <typed list (variable)> )
   <action-def body>)
<action functor> ::= <name>
<action-def body> ::=
  [:vars (<typed list(variable)>)] :existential-preconditions :conditional-effects
  [:precondition <GD>]
  [:expansion <action spec>] :action-expansions
  [:expansion :methods] :action-expansions
  [:maintain <GD>] :action-expansions
  [:effect <effect>]
  [:only-in-expansions <boolean>] :action-expansions
```

PDDL Goal Descriptions

```
<GD> ::= <atomic formula(term)>
<GD> ::= (and <GD>+)
<GD> ::= <literal(term)>
<GD> ::= :disjunctive-preconditions (or <GD>+)
<GD> ::= :disjunctive-preconditions (not <GD>)
<GD> ::= :disjunctive-preconditions (imply <GD> <GD>)
<GD> ::= :existential-preconditions (exists (<typed list(variable)>) <GD> )
<GD> ::= :universal-preconditions (forall (<typed list(variable)>) <GD> )
<literal(t)> ::= <atomic formula(t)>
<literal(t)> ::= (not <atomic formula(t)>)
<atomic formula(t)> ::= (<predicate> t*)
<term> ::= <name>
```

PDDL Effects

```
<effect> ::= (and <effect>+)  
<effect> ::= <atomic formula(term)>  
<effect> ::= (not <atomic formula(term)>)  
<effect> ::=:conditional-effects  
    (forall (<variable>*) <effect>)  
<effect> ::=:conditional-effects  
    (when <GD> <effect>)  
<effect> ::=:fluents (change <fluent> <expression>)
```

Example: DWR Action

```
:: moves a robot between two adjacent locations  
(:action move  
 :parameters (?r - robot ?from ?to - location)  
 :precondition (and  
    (adjacent ?from ?to) (at ?r ?from)  
    (not (occupied ?to)))  
 :effect (and  
    (at ?r ?to) (occupied ?to)  
    (not (occupied ?from)) (not (at ?r ?from)) ))
```

PDDL Problem Descriptions

```

<problem> ::= (define (problem <name>)
  (:domain <name>)
  [<require-def>]
  [<situation> ]
  [<object declaration> ]
  [<init>]
  <goal>+
  [<length-spec> ])
<object declaration> ::= (:objects <typed list (name)>)
<situation> ::= (:situation <initsit name>)
<initsit name> ::= <name>
<init> ::= (:init <literal(name)>+)
<goal> ::= (:goal <GD>)
<goal> ::= :action-expansions (:expansion <action spec(action-term)>)
<length-spec> ::= (:length [(:serial <integer>)] [(:parallel <integer>)])
  
```

Example: DWR Problem

```

;; a simple DWR problem with 1 robot and 2
locations
(define (problem dwrpb1)
  (:domain dock-worker-robot)
  (:objects
    r1 - robot
    l1 l2 - location
    k1 k2 - crane
    p1 q1 p2 q2 - pile
    ca cb cc cd ce cf pallet - container)
  (:init
    (adjacent l1 l2)
    (adjacent l2 l1)
    (attached p1 l1)
    (attached q1 l1)
    (attached p2 l2)
    (attached q2 l2)
    (belong k1 l1)
    (belong k2 l2)
    (in ca p1) (in cb p1) (in cc p1)
    (on ca pallet) (on cb ca) (on cc cb)
    (top cc p1)
    (in cd q1) (in ce q1) (in cf q1)
    (on cd pallet) (on ce cd) (on cf ce)
    (top cf q1)
    (top pallet p2)
    (top pallet q2)
    (at r1 l1)
    (unloaded r1)
    (occupied l1)
    (empty k1)
    (empty k2))
  ;; task is to move all containers to locations l2
  ;; ca and cc in pile p2, the rest in q2
  (:goal (and
    (in ca p2) (in cc p2)
    (in cb q2) (in cd q2) (in ce q2) (in cf q2))))
  
```

Overview


- The STRIPS Representation
- The Planning Domain Definition Language (PDDL)
- **Problem-Solving by Search**
- Heuristic Search
- Forward State-Space Search
- Backward State-Space Search
- The STRIPS Planner

Search Problems

- initial state
- set of possible actions/applicability conditions
 - successor function: $state \rightarrow \text{set of } \langle action, state \rangle$
 - successor function + initial state = state space
 - path (solution)
- goal
 - goal state or goal test function
- path cost function
 - for optimality
 - assumption: path cost = sum of step costs

Missionaries and Cannibals: Initial State and Actions


- initial state:
 - all missionaries, all cannibals, and the boat are on the left bank
- 5 possible actions:
 - one missionary crossing
 - one cannibal crossing
 - two missionaries crossing
 - two cannibals crossing
 - one missionary and one cannibal crossing



Missionaries and Cannibals: Successor Function

<i>state</i>	set of <i><action, state></i>
(L:3m,3c,b-R:0m,0c) →	{<2c, (L:3m,1c-R:0m,2c,b)>, <1m1c, (L:2m,2c-R:1m,1c,b)>, <1c, (L:3m,2c-R:0m,1c,b)>}
(L:3m,1c-R:0m,2c,b) →	{<2c, (L:3m,3c,b-R:0m,0c)>, <1c, (L:3m,2c,b-R:0m,1c)>}
(L:2m,2c-R:1m,1c,b) →	{<1m1c, (L:3m,3c,b-R:0m,0c)>, <1m, (L:3m,2c,b-R:0m,1c)>}
⋮	⋮

Missionaries and Cannibals: State Space




State-Space Search and the STRIPS Planner

35

Missionaries and Cannibals: Goal State and Path Cost

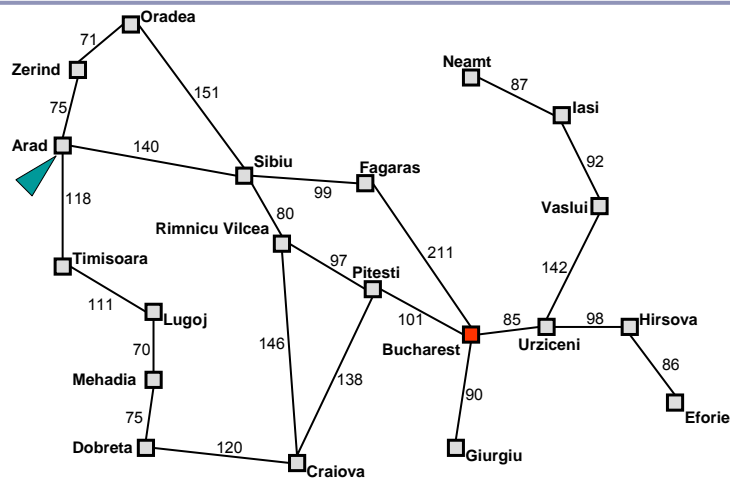
- goal state:
 - all missionaries, all cannibals, and the boat are on the right bank
- path cost
 - step cost: 1 for each crossing
 - path cost: number of crossings = length of path
- solution path:
 - 4 optimal solutions
 - cost: 11



State-Space Search and the STRIPS Planner

36

Real-World Problem: Touring in Romania



State-Space Search and the STRIPS Planner

37

Touring Romania: Search Problem Definition

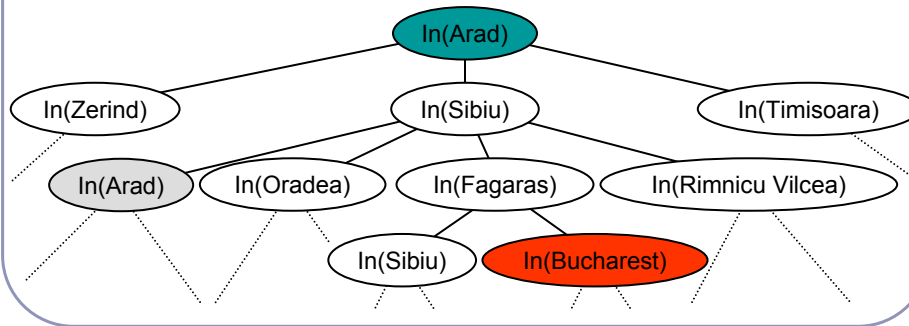
- initial state:
 - In(Arad)
- possible Actions:
 - DriveTo(Zerind), DriveTo(Sibiu), DriveTo(Timisoara), etc.
- goal state:
 - In(Bucharest)
- step cost:
 - distances between cities

State-Space Search and the STRIPS Planner

38

Search Trees

- **search tree**: tree structure defined by initial state and successor function
- Touring Romania (partial search tree):



State-Space Search and the STRIPS Planner

39

Search Nodes

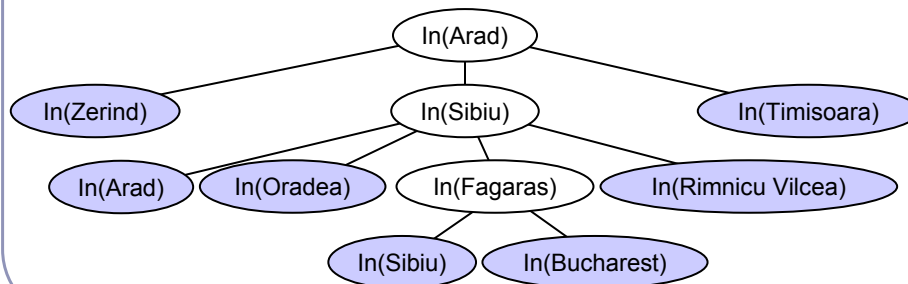
- **search nodes**: the nodes in the search tree
- **data structure**:
 - *state*: a state in the state space
 - *parent node*: the immediate predecessor in the search tree
 - *action*: the action that, performed in the parent node's state, leads to this node's state
 - *path cost*: the total cost of the path leading to this node
 - *depth*: the depth of this node in the search tree

State-Space Search and the STRIPS Planner

40

Fringe Nodes in Touring Romania Example

fringe nodes: nodes that have not been expanded



State-Space Search and the STRIPS Planner

41

Search (Control) Strategy

- search or control strategy: an effective method for scheduling the application of the successor function to expand nodes
 - selects the next node to be expanded from the fringe
 - determines the order in which nodes are expanded
 - aim: produce a goal state as quickly as possible
- examples:
 - LIFO/FIFO-queue for fringe nodes
 - alphabetical ordering

State-Space Search and the STRIPS Planner

42

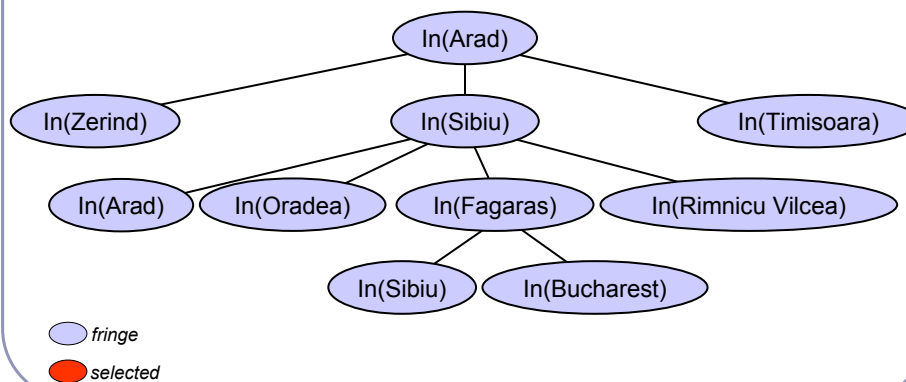
General Tree Search Algorithm

```
function treeSearch(problem, strategy)  
  fringe ← { new  
    searchNode(problem.initialState) }  
  loop  
    if empty(fringe) then return failure  
    node ← selectFrom(fringe, strategy)  
    if problem.goalTest(node.state) then  
      return pathTo(node)  
    fringe ← fringe + expand(problem, node)
```

State-Space Search and the STRIPS Planner

43

General Search Algorithm: Touring Romania Example




State-Space Search and the STRIPS Planner

44

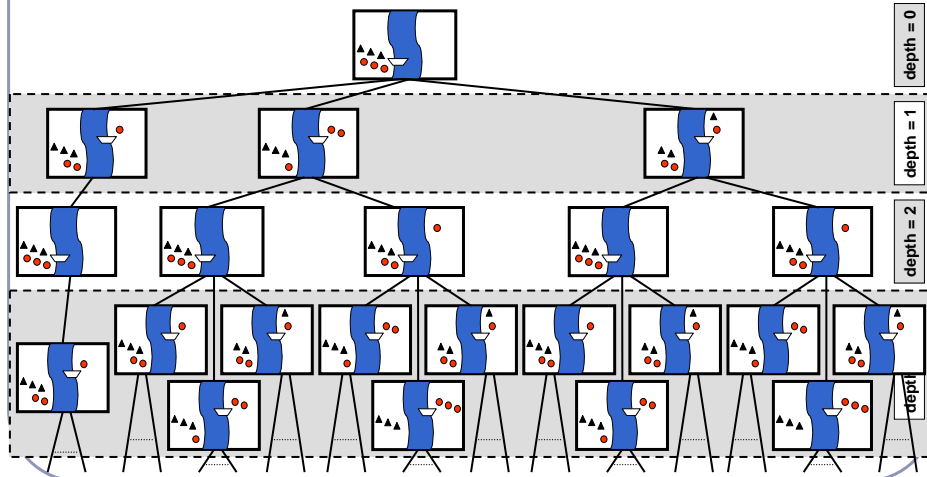
Uninformed vs. Informed Search

- uninformed search (blind search)
 - no additional information about states beyond problem definition
 - only goal states and non-goal states can be distinguished
- informed search (heuristic search)
 - additional information about how “promising” a state is available

Breadth-First Search: Missionaries and Cannibals



Depth-First Search: Missionaries and Cannibals



State-Space Search and the STRIPS Planner

47

Iterative Deepening Search


- *strategy*:
 - based on depth-limited (depth-first) search
 - repeat search with gradually increasing depth limit until a goal state is found
- *implementation*:
for $depth \leftarrow 0$ **to** ∞ **do**
 $result \leftarrow \text{depthLimitedSearch}(\text{problem}, \text{depth})$
 if $result \neq \text{cutoff}$ **then return** $result$

State-Space Search and the STRIPS Planner

48

Discovering Repeated States: Potential Savings

- sometimes repeated states are unavoidable, resulting in infinite search trees
- checking for repeated states:
 - infinite search tree \Rightarrow finite search tree
 - finite search tree \Rightarrow exponential reduction



State-Space Search and the STRIPS Planner

49

Overview

- The STRIPS Representation
- The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- **Heuristic Search**
- Forward State-Space Search
- Backward State-Space Search
- The STRIPS Planner

State-Space Search and the STRIPS Planner

50

Best-First Search

- an instance of the general tree search or graph search algorithm
 - strategy: select next node based on an evaluation function $f: \text{state space} \rightarrow \mathbb{R}$
 - select node with lowest value $f(n)$
- implementation:
selectFrom(*fringe*, *strategy*)
 - priority queue: maintains fringe in ascending order of f -values

Heuristic Functions

- heuristic function $h: \text{state space} \rightarrow \mathbb{R}$
- $h(n)$ = estimated cost of the cheapest path from node n to a goal node
- if n is a goal node then $h(n)$ must be 0
- heuristic function encodes problem-specific knowledge in a problem-independent way

Greedy Best-First Search

- use heuristic function as evaluation function: $f(n) = h(n)$
 - always expands the node that is closest to the goal node
 - eats the largest chunk out of the remaining distance, hence, “greedy”


Touring in Romania: Heuristic

- $h_{SLD}(n)$ = straight-line distance to Bucharest

Arad	366	Hirsova	151	Rimnicu	193
Bucharest	0	Iasi	226	Vilcea	
Craiova	160	Lugoj	244	Sibiu	253
Dobreta	242	Mehadia	241	Timisoara	329
Eforie	161	Neamt	234	Urziceni	80
Fagaras	176	Oradea	380	Vaslui	199
Giurgiu	77	Pitesti	100	Zerind	374

Greediness

- greediness is susceptible to false starts



- repeated states may lead to infinite oscillation



State-Space Search and the STRIPS Planner

55

A* Search

- best-first search where

$$f(n) = h(n) + g(n)$$

- $h(n)$ the heuristic function (as before)
- $g(n)$ the cost to reach the node n

- evaluation function:

$f(n)$ = estimated cost of the cheapest solution through n

- A* search is optimal if $h(n)$ is admissible

State-Space Search and the STRIPS Planner

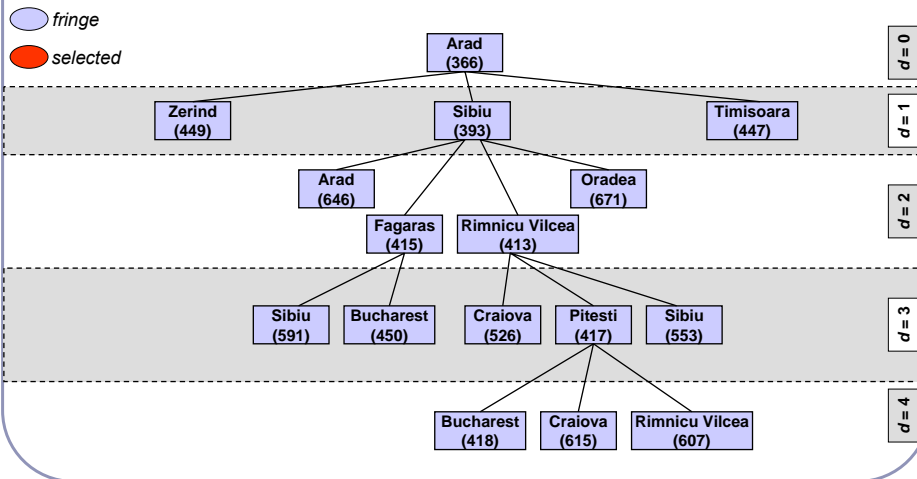
56

Admissible Heuristics

A heuristic $h(n)$ is admissible if it *never overestimates* the distance from n to the nearest goal node.

- example: h_{SLD}
- A* search: If $h(n)$ is admissible then $f(n)$ never overestimates the true cost of a solution through n .

A* (Best-First) Search: Touring Romania



Optimality of A* (Tree Search)

Theorem:

A* using tree search is optimal if the heuristic $h(n)$ is admissible.

A*: Optimally Efficient

- A* is optimally efficient for a given heuristic function:
no other optimal algorithm is guaranteed to expand fewer nodes than A*.
- any algorithm that does not expand all nodes with $f(n) < C^*$ runs the risk of missing the optimal solution

A* and Exponential Space

- A* has worst case time and space complexity of $O(b^l)$
- exponential growth of the fringe is normal
 - exponential time complexity may be acceptable
 - exponential space complexity will exhaust any computer's resources all too quickly

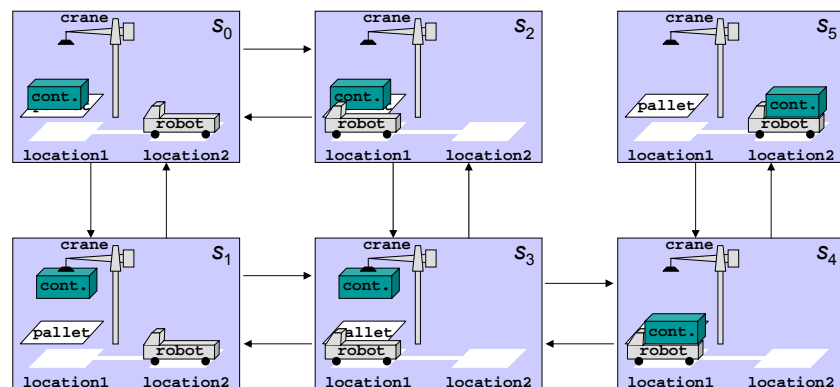
Overview

- The STRIPS Representation
- The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- Heuristic Search
- **Forward State-Space Search**
- Backward State-Space Search
- The STRIPS Planner

State-Space Search

- idea: apply standard search algorithms (breadth-first, depth-first, A*, etc.) to planning problem:
 - search space is subset of state space
 - nodes correspond to world states
 - arcs correspond to state transitions
 - path in the search space corresponds to plan

DWR Example: State Space



Search Problems

- initial state
- set of possible actions/applicability conditions
 - successor function: $state \rightarrow$ set of $\langle action, state \rangle$
 - successor function + initial state = state space
 - path (solution)
- goal
 - goal state or goal test function
- path cost function
 - for optimality
 - assumption: path cost = sum of step costs

State-Space Planning as a Search Problem

- given: statement of a planning problem
 $P=(O,s_i,g)$
- define the search problem as follows:
 - initial state: s_i
 - goal test for state s : s satisfies g
 - path cost function for plan π : $|\pi|$
 - successor function for state s : $\Gamma(s)$

Reachable Successor States

- The successor function $\Gamma^m: 2^S \rightarrow 2^S$ for a STRIPS domain $\Sigma = (S, A, \gamma)$ is defined as:
 - $\Gamma(s) = \{\gamma(s, a) \mid a \in A \text{ and } a \text{ applicable in } s\}$ for $s \in S$
 - $\Gamma(\{s_1, \dots, s_n\}) = \bigcup_{k \in [1, n]} \Gamma(s_k)$
 - $\Gamma^0(\{s_1, \dots, s_n\}) = \{s_1, \dots, s_n\}$
 - $\Gamma^m(\{s_1, \dots, s_n\}) = \Gamma(\Gamma^{m-1}(\{s_1, \dots, s_n\}))$
- The transitive closure of Γ defines the set of all reachable states:
 - $\Gamma^>(s) = \bigcup_{k \in [0, \infty)} \Gamma^k(\{s\})$ for $s \in S$

Solution Existence

- **Proposition:** A STRIPS planning problem $\mathcal{P} = (\Sigma, s_i, g)$ (and a statement of such a problem $P = (O, s_i, g)$) has a solution iff $S_g \cap \Gamma^>(\{s_i\}) \neq \{\}$.

Forward State-Space Search Algorithm

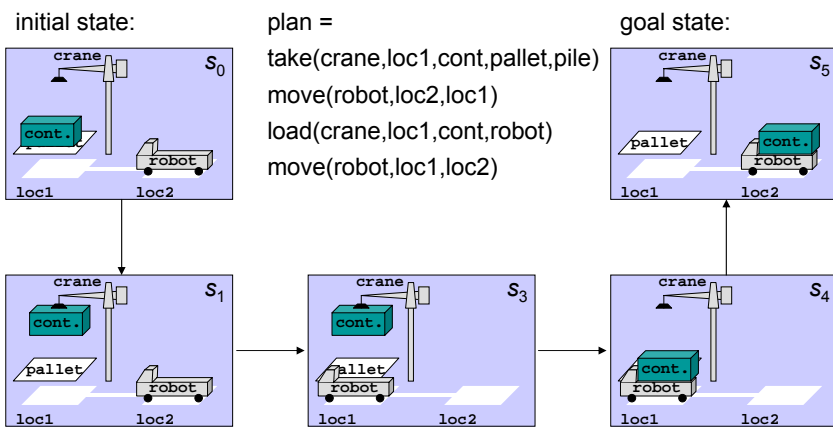
```

function fwdSearch( $O, s_i, g$ )
   $state \leftarrow s_i$ 
   $plan \leftarrow \langle \rangle$ 
  loop
    if  $state.satisfies(g)$  then return  $plan$ 
     $applicables \leftarrow$ 
      {ground instances from  $O$  applicable in  $state$ }
    if  $applicables.isEmpty()$  then return failure
     $action \leftarrow applicables.chooseOne()$ 
     $state \leftarrow \gamma(state, action)$ 
     $plan \leftarrow plan \cdot \langle action \rangle$ 
  
```

State-Space Search and the STRIPS Planner

69

DWR Example: Forward Search



State-Space Search and the STRIPS Planner

70

Finding Applicable Actions: Algorithm

```
function addApplicables(A, op, precs,  $\sigma$ , s)
  if precs+.isEmpty() then
    for every np in precs- do
      if s.falsifies( $\sigma$ (np)) then return
      A.add( $\sigma$ (op))
  else
    pp  $\leftarrow$  precs+.chooseOne()
    for every sp in s do
       $\sigma'$   $\leftarrow$   $\sigma$ .extend(sp, pp)
      if  $\sigma'$ .isValid() then
        addApplicables(A, op, (precs - pp),  $\sigma'$ , s)
```

Properties of Forward Search

- **Proposition:** fwdSearch is sound, i.e. if the function returns a plan as a solution then this plan is indeed a solution.
 - proof idea: show (by induction) $state = \gamma(s_i, plan)$ at the beginning of each iteration of the loop
- **Proposition:** fwdSearch is complete, i.e. if there exists solution plan then there is an execution trace of the function that will return this solution plan.
 - proof idea: show (by induction) there is an execution trace for which *plan* is a prefix of the sought plan

Making Forward Search Deterministic

- idea: use depth-first search
 - problem: infinite branches
 - solution: prune repeated states
- pruning: cutting off search below certain nodes
 - safe pruning: guaranteed not to prune every solution
 - strongly safe pruning: guaranteed not to prune every optimal solution
 - example: prune below nodes that have a predecessor that is an equal state (no repeated states)

Overview

- The STRIPS Representation
- The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- Heuristic Search
- Forward State-Space Search
- Backward State-Space Search
- The STRIPS Planner

The Problem with Forward Search

- number of actions applicable in any given state is usually very large
- branching factor is very large
- forward search for plans with more than a few steps not feasible

- idea: search backwards from the goal
- problem: many goal states

Relevance and Regression Sets

- Let $\mathcal{P}=(\Sigma, s_i, g)$ be a STRIPS planning problem. An action $a \in A$ is relevant for g if
 - $g \cap \text{effects}(a) \neq \{\}$ and
 - $g^+ \cap \text{effects}^-(a) = \{\}$ and $g^- \cap \text{effects}^+(a) = \{\}$.
- The regression set of g for a relevant action $a \in A$ is:
 - $\gamma^{-1}(g, a) = (g - \text{effects}(a)) \cup \text{precond}(a)$

Regression Function

- The regression function Γ^{-m} for a STRIPS domain $\Sigma=(S,A,\gamma)$ on L is defined as:
 - $\Gamma^{-1}(g)=\{g^{-1}(g,a) \mid a \in A \text{ is relevant for } g\}$ for $g \in 2^L$
 - $\Gamma^0(\{g_1, \dots, g_n\}) = \{g_1, \dots, g_n\}$
 - $\Gamma^{-1}(\{g_1, \dots, g_n\}) = \bigcup_{(k \in [1, n])} \Gamma^{-1}(g_k)$
 - $\Gamma^{-m}(\{g_1, \dots, g_n\}) = \Gamma^{-1}(\Gamma^{-(m-1)}(\{g_1, \dots, g_n\}))$
- $\left. \begin{array}{l} \bullet \\ \bullet \\ \bullet \\ \bullet \end{array} \right\} g_1, \dots, g_n \in 2^L$
- The transitive closure of Γ^{-1} defines the set of all regression sets:
 - $\Gamma^<(g) = \bigcup_{(k \in [0, \infty])} \Gamma^{-k}(\{g\})$ for $g \in 2^L$

State-Space Planning as a Search Problem

- given: statement of a planning problem $P=(O, s_i, g)$
- define the search problem as follows:
 - initial search state: g
 - goal test for state s : s_i satisfies s
 - path cost function for plan π : $|\pi|$
 - successor function for state s : $\Gamma^{-1}(s)$

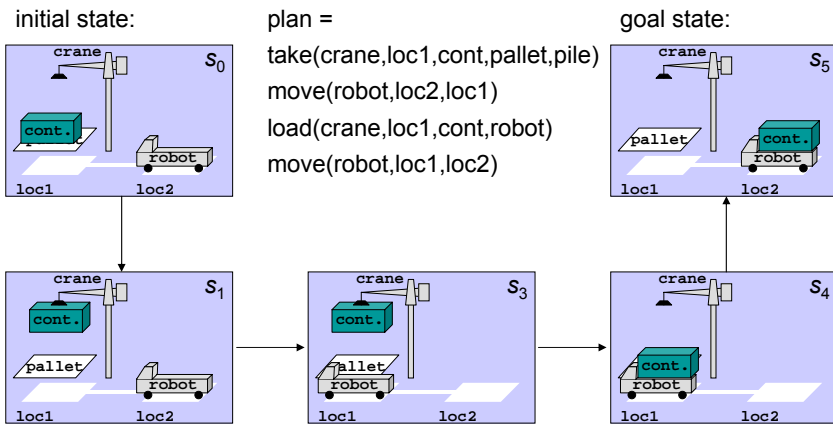
Solution Existence

- **Proposition:** A propositional planning problem $\mathcal{P}=(\Sigma, s_i, g)$ (and a statement of such a problem $P=(O, s_i, g)$) has a solution iff $\exists s \in \Gamma^<(\{g\}) : s_i$ satisfies s .

Ground Backward State-Space Search Algorithm

```
function groundBwdSearch( $O, s_i, g$ )  
   $subgoal \leftarrow g$   
   $plan \leftarrow \langle \rangle$   
  loop  
    if  $s_i$ .satisfies( $subgoal$ ) then return  $plan$   
     $applicables \leftarrow$   
      {ground instances from  $O$  relevant for  $subgoal$ }  
    if  $applicables$ .isEmpty() then return failure  
     $action \leftarrow applicables$ .chooseOne()  
     $subgoal \leftarrow \gamma^{-1}(subgoal, action)$   
     $plan \leftarrow \langle action \rangle \cdot plan$ 
```


DWR Example: Backward Search



State-Space Search and the STRIPS Planner

81

Example: Regression with Operators

- goal: $at(robot,loc1)$
- operator: $move(r,l,m)$
 - precondition: $adjacent(l,m), at(r,l), \neg occupied(m)$
 - effects: $at(r,m), occupied(m), \neg occupied(l), \neg at(r,l)$
- actions: $move(robot,l,loc1)$
 - $l=?$
 - many options increase branching factor
- lifted backward search: use partially instantiated operators instead of actions

State-Space Search and the STRIPS Planner

82

Lifted Backward State-Space Search Algorithm

```


function liftedBwdSearch( $O, s_i, g$ )
   $subgoal \leftarrow g$ 
   $plan \leftarrow \langle \rangle$ 
  loop
    if  $\exists \sigma: s_i$ .satisfies( $\sigma(subgoal)$ ) then return  $\sigma(plan)$ 
     $applicables \leftarrow$ 
       $\{(o, \sigma) \mid o \in O \text{ and } \sigma(o) \text{ relevant for } subgoal\}$ 
    if  $applicables.isEmpty()$  then return failure
     $action \leftarrow applicables.chooseOne()$ 
     $subgoal \leftarrow \gamma^{-1}(\sigma(subgoal), \sigma(o))$ 
     $plan \leftarrow \sigma(\langle action \rangle) \cdot \sigma(plan)$ 

```

State-Space Search and the STRIPS Planner

83

DWR Example: Lifted Backward Search



- **initial state:** $s_0 = \{attached(pile, loc1), in(cont, pile), top(cont, pile), on(cont, pallet), belong(crane, loc1), empty(crane), adjacent(loc1, loc2), adjacent(loc2, loc1), at(robot, loc2), occupied(loc2), unloaded(robot)\}$
- **operator:** $move(r, l, m)$
 - **precond:** $adjacent(l, m), at(r, l), \neg occupied(m)$
 - **effects:** $at(r, m), occupied(m), \neg occupied(l), \neg at(r, l)$
- **liftedBwdSearch**($\{move(r, l, m)\}, s_0, \{at(robot, loc1)\}$)
- $\exists \sigma: s_i$.satisfies($\sigma(subgoal)$): no
- **applicables** = $\{(move(r, l_1, m_1), \{r_1 \leftarrow robot, m_1 \leftarrow loc1\})\}$
- **subgoal** = $\{adjacent(l_1, loc1), at(robot, l_1), \neg occupied(loc1)\}$
- **plan** = $\langle move(robot, l_1, loc1) \rangle$
- $\exists \sigma: s_i$.satisfies($\sigma(subgoal)$): yes
 $\sigma = \{l_1 \leftarrow loc1\}$

State-Space Search and the STRIPS Planner

84

Properties of Backward Search

- **Proposition:** liftedBwdSearch is sound, i.e. if the function returns a plan as a solution then this plan is indeed a solution.
 - proof idea: show (by induction) $subgoal = \gamma^{-1}(g, plan)$ at the beginning of each iteration of the loop
- **Proposition:** liftedBwdSearch is complete, i.e. if there exists solution plan then there is an execution trace of the function that will return this solution plan.
 - proof idea: show (by induction) there is an execution trace for which $plan$ is a suffix of the sought plan

Avoiding Repeated States

- search space:
 - let g_i and g_k be sub-goals where g_i is an ancestor of g_k in the search tree
 - let σ be a substitution such that $\sigma(g_i) \subseteq g_k$
- pruning:
 - then we can prune all nodes below g_k

Overview

- The STRIPS Representation
- The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- Heuristic Search
- Forward State-Space Search
- Backward State-Space Search
- **The STRIPS Planner**

Problems with Backward Search

- state space still too large to search efficiently
- STRIPS idea:
 - only work on preconditions of the last operator added to the plan
 - if the current state satisfies all of an operator's preconditions, commit to this operator

Ground-STRIPS Algorithm


```
function groundStrips( $O,s,g$ )  
  plan  $\leftarrow \langle \rangle$   
  loop  
    if  $s$ .satisfies( $g$ ) then return plan  
    applicables  $\leftarrow$   
      {ground instances from  $O$  relevant for  $g$ - $s$ }  
    if applicables.isEmpty() then return failure  
    action  $\leftarrow$  applicables.chooseOne()  
    subplan  $\leftarrow$  groundStrips( $O,s$ ,action.preconditions())  
    if subplan = failure then return failure  
     $s \leftarrow \gamma(s, \text{subplan} \cdot \langle \text{action} \rangle)$   
    plan  $\leftarrow$  plan  $\cdot$  subplan  $\cdot \langle \text{action} \rangle$ 
```

State-Space Search and the STRIPS Planner

89

Problems with STRIPS

- STRIPS is incomplete:
 - cannot find solution for some problems, e.g. interchanging the values of two variables
 - cannot find optimal solution for others, e.g. Sussman anomaly:






State-Space Search and the STRIPS Planner

90

STRIPS and the Sussman Anomaly (1)

- achieve on(A,B)
 - put C from A onto table
 - put A onto B
- achieve on(B,C)
 - put A from B onto table
 - put B onto C
- re-achieve on(A,B)
 - put A onto B







State-Space Search and the STRIPS Planner

91

STRIPS and the Sussman Anomaly (2)

- achieve on(B,C)
 - put B onto C
- achieve on(A,B)
 - put B from C onto table
 - put C from A onto table
 - put A onto B
- re-achieve on(B,C)
 - put A from B onto table
 - put B onto C
- re-achieve on(A,B)
 - put A onto B



State-Space Search and the STRIPS Planner

92

Interleaving Plans for an Optimal Solution

- shortest solution achieving $on(A,B)$:
- shortest solution for $on(A,B)$ and $on(B,C)$:

put C from A onto table

put A onto B

- shortest solution achieving $on(B,C)$:

put B onto C

Overview

- The STRIPS Representation
- The Planning Domain Definition Language (PDDL)
- Problem-Solving by Search
- Heuristic Search
- Forward State-Space Search
- Backward State-Space Search
- **The STRIPS Planner**