

Set 9: Planning; STRIPs Planning Systems

ICS 271 Fall 2012

Outline: Planning


- Classical Planning:
 - Situation calculus
 - PDDL: Planning domain definition language
- STRIPS Planning
- Planning graphs
- Readings: Russel and Norvig chapter 10

The Situation Calculus

- A **goal** can be described by a sentence:
if we want to have a block on B $(\exists x)On(x, B)$
- **Planning**: finding a set of actions to achieve a goal sentence.
- **Situation Calculus** (McCarthy, Hayes, 1969, Green 1969)
 - A Predicate Calculus formalization of *states*, *actions*, and their *effects*.
 - S_0 state in the figure can be described by:

$On(B, A) \wedge On(A, C) \wedge On(C, F_1) \wedge Clear(B) \wedge clear(F_1)$

we reify the state and
include them as arguments



On (B , A)
On (A , C)
On (C , F1)
Clear (B)
Clear (F1)

The Situation Calculus (continued)

- The atoms denotes relations over states called **fluents**.

- We can also have. $O_n(B, A, S_0) \wedge O_n(A, C, S_0) \wedge O_n(C, F_1, S_0) \wedge clear(B, S_0)$

$$(\forall x, y, s)[On(x, y, s) \wedge \neg(y = F_1) \rightarrow \neg Clear(y, s)]$$

$$(\forall s)Clear(F_1, s)$$

- Knowledge about state and actions = predicate calculus theory.
- Inference can be used to answer:
 - Is there a state satisfying a goal?
 - How can the present state be transformed into that state by actions?
The answer is a *plan*

Representing Actions

- Reify the actions: denote an action by a symbol
- actions are **functions**
- $\text{move}(B,A,F1)$: move block A from block B to F1
- $\text{move}(x,y,z)$ - action schema
- **do**: A function constant, **do** denotes a function that maps actions and states into states

$$- \quad \text{do}(\alpha, \sigma) \rightarrow \sigma_1$$

The diagram illustrates the function $\text{do}(\alpha, \sigma) \rightarrow \sigma_1$. The word "action" has an arrow pointing to the symbol α . The word "state" has an arrow pointing to the symbol σ . The symbol σ_1 is the result of the function.

Representing Actions (continued)

- Express the effects of actions.
 - Example: (on, move) (expresses the effect of move on “On”)

$[On(x, y, s) \wedge Clear(x, s) \wedge Clear(z, s) \wedge (x \neq z)$

$\rightarrow On(x, z, do(move(x, y, z), s))]$

negative :

$[On(x, y, s) \wedge Clear(x, s) \wedge Clear(z, s) \wedge (x \neq z)$

$\rightarrow \neg On(x, y, do(move(x, y, z), s))]$

- Positive: describes how action makes a fluent true
- Negative : describes how action makes a fluent false
- Antecedent: pre-condition for actions
- Consequent: how the fluent is changed

Frame Axioms

- Not everything true can be inferred
On(C,F1) remains true but cannot be inferred
- Actions have local effect
 - We need **frame axioms** for each action and each fluent that does not change as a result of the action
 - example: frame axioms for (move, on)
 - If a block is on another block and *move* is not relevant, it will stay the same.

- Positive:

$$[On(x, y, s) \wedge (x \neq u)] \rightarrow On(x, y, do(move(u, v, z), s))$$

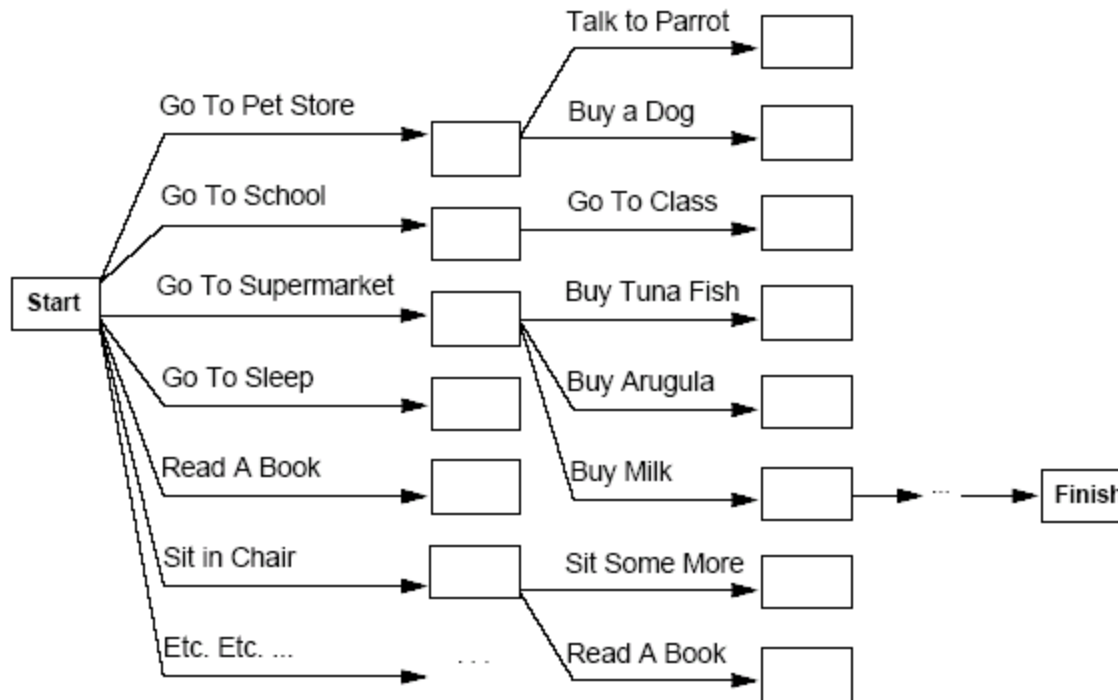
- negative

$$(\neg On(x, y, s) \wedge [(x \neq u) \vee (y \neq z)]) \rightarrow \neg On(x, y, do(move(u, v, z), s))$$

STRIPS Planning systems
PDDL: Planning Domain Definition
Language

Search vs. planning

Consider the task *get milk, bananas, and a cordless drill*
Standard search algorithms seem to fail miserably:




After-the-fact heuristic/goal test inadequate

STRIPS: describing goals and state

Factored representation of states

- On(B,A)
- On(A,C)
- On(C,F1)
- Clear(B)
- Clear(F1)



```
On (B, A)
On (A, C)
On (C, F1)
Clear (B)
Clear (F1)
```

- The formula describes a set of world states
- Planning search for a formula satisfying a goal description
- **State descriptions:** conjunctions of ground literals.
- Also universal formulas: $\text{On}(x,y) \rightarrow (y=F1)$ or $\sim\text{Clear}(y)$
- Goal wff: $\exists x.g(x) \wedge f(y)$
- Given a goal wff, the search algorithm looks for a sequence of **actions**

That transform into a state description that entails the goal wff.

STRIPS Description of Operators

- A STRIPS operator has 3 parts:
 - A set, PC (*preconditions*) of ground literals
 - A set D, of ground literals called the *delete list*
 - A set A, of ground literals called *add list*
- Usually described by Schema: *Move(x,y,z)*
 - PC: *On(x,y) and Clear(x) and Clear(z)*
 - D: *Clear(z) , On(x,y)*
 - A: *On(x,z), Clear(y), Clear(F1)*
- A state S1 is created applying operator O by adding A and deleting D from S1.

STRIPS operators

Tidily arranged actions descriptions, restricted language

ACTION: $Buy(x)$

PRECONDITION: $At(p), Sells(p, x)$

EFFECT: $Have(x)$

[Note: this abstracts away many important details!]

Restricted language \Rightarrow efficient algorithm

Precondition: conjunction of positive literals

Effect: conjunction of literals


A complete set of STRIPS operators can be translated into a set of successor-state axioms

$At(p) Sells(p, x)$

Buy(x)

$Have(x)$

Example: the move operator




Precondition:

On (B, A)

Clear (B)

Clear (F1)



PDDL

- A language that yields a search problem
- A state is a set of ground literals
- Closed world assumption: fluent that are not mentioned are false.
- Action schema:

Action(Fly(p,from,to),


Precond: At(p,from) & Plane(p) & Airport(from) & Airport(to)

Effect: not At(p,from) & At(p,to))

- *The schema consists of precondition and effect lists*
- *PDDL is very close to STRIP language*

- *A set of action schemas is a definition of a planning domain.*
- *A specific problem is defined by an initial state (a set of ground literals) and a goal: conjunction of literals, some not grounded (At(p,SFO), Plane(p))*

The block world



```
Init(On(A, Table)  $\wedge$  On(B, Table)  $\wedge$  On(C, Table)
   $\wedge$  Block(A)  $\wedge$  Block(B)  $\wedge$  Block(C)
   $\wedge$  Clear(A)  $\wedge$  Clear(B)  $\wedge$  Clear(C))
Goal(On(A, B)  $\wedge$  On(B, C))
Action(Move(b, x, y),
  PRECOND: On(b, x)  $\wedge$  Clear(b)  $\wedge$  Clear(y)  $\wedge$  Block(b)  $\wedge$ 
    {b  $\neq$  x}  $\wedge$  {b  $\neq$  y}  $\wedge$  {x  $\neq$  y},
  EFFECT: On(b, y)  $\wedge$  Clear(x)  $\wedge$   $\neg$  On(b, x)  $\wedge$   $\neg$  Clear(y))
Action(MoveToTable(b, x),
  PRECOND: On(b, x)  $\wedge$  Clear(b)  $\wedge$  Block(b)  $\wedge$  {b  $\neq$  x},
  EFFECT: On(b, Table)  $\wedge$  Clear(x)  $\wedge$   $\neg$  On(b, x))
```

Figure 11.4 A planning problem in the blocks world: building a three-block tower. One solution is the sequence [*Move*(B, Table, C), *Move*(A, Table, B)].

A STRIP/PDDL description of an aircargo transportation problem

Problem: flying cargo in planes from one location to another

```
Init{At(C1, SFO) ∧ At(C2, JFK) ∧ At(P1, SFO) ∧ At(P2, JFK)
  ∧ Cargo(C1) ∧ Cargo(C2) ∧ Plane(P1) ∧ Plane(P2)
  ∧ Airport(JFK) ∧ Airport(SFO))
Goal{At(C1, JFK) ∧ At(C2, SFO)}
Action{Load(c, p, a),
  PRECOND: At(c, a) ∧ At(p, a) ∧ Cargo(c) ∧ Plane(p) ∧ Airport(a)
  EFFECT: ¬ At(c, a) ∧ In(c, p))
Action{Unload(c, p, a),
  PRECOND: In(c, p) ∧ At(p, a) ∧ Cargo(c) ∧ Plane(p) ∧ Airport(a)
  EFFECT: At(c, a) ∧ ¬ In(c, p))
Action{Fly(p, from, to),
  PRECOND: At(p, from) ∧ Plane(p) ∧ Airport(from) ∧ Airport(to)
  EFFECT: ¬ At(p, from) ∧ At(p, to)}
```

Figure 11.2 A STRIPS problem involving transportation of air cargo between airports.

In(*c*,*p*)- cargo *c* is inside plane *p*

At(*x*,*a*) – object *x* is at airport *a*

STRIP for spare tire problem

Problem: Changing a flat tire

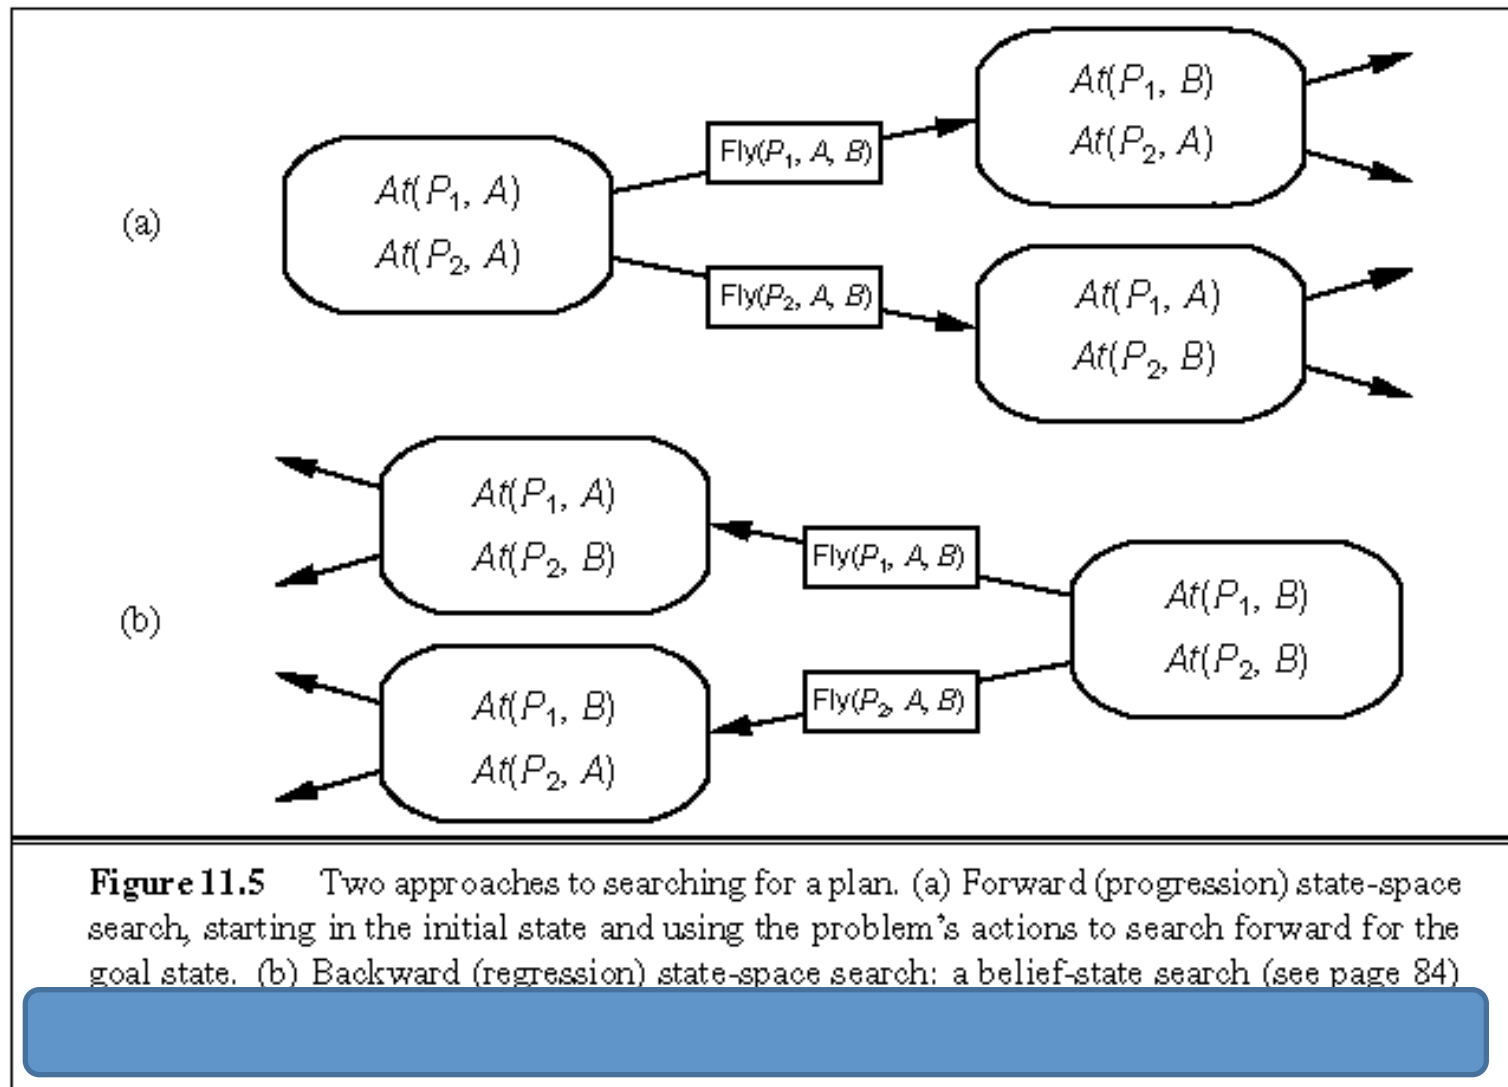
```
Init(At{Flat, Axle}  $\wedge$  At{Spare, Trunk})  
Goal(At{Spare, Axle})  
Action{Remove{Spare, Trunk},  
  PRECOND: At{Spare, Trunk}  
  EFFECT:  $\neg$  At{Spare, Trunk}  $\wedge$  At{Spare, Ground})  
Action{Remove{Flat, Axle},  
  PRECOND: At{Flat, Axle}  
  EFFECT:  $\neg$  At{Flat, Axle}  $\wedge$  At{Flat, Ground})  
Action{PutOn{Spare, Axle},  
  PRECOND: At{Spare, Ground}  $\wedge$   $\neg$  At{Flat, Axle}  
  EFFECT:  $\neg$  At{Spare, Ground}  $\wedge$  At{Spare, Axle})  
Action{LeaveOvernight,  
  PRECOND:  
  EFFECT:  $\neg$  At{Spare, Ground}  $\wedge$   $\neg$  At{Spare, Axle}  $\wedge$   $\neg$  At{Spare, Trunk}  
           $\wedge$   $\neg$  At{Flat, Ground}  $\wedge$   $\neg$  At{Flat, Axle})
```

Figure 11.3 The simple spare tire problem.

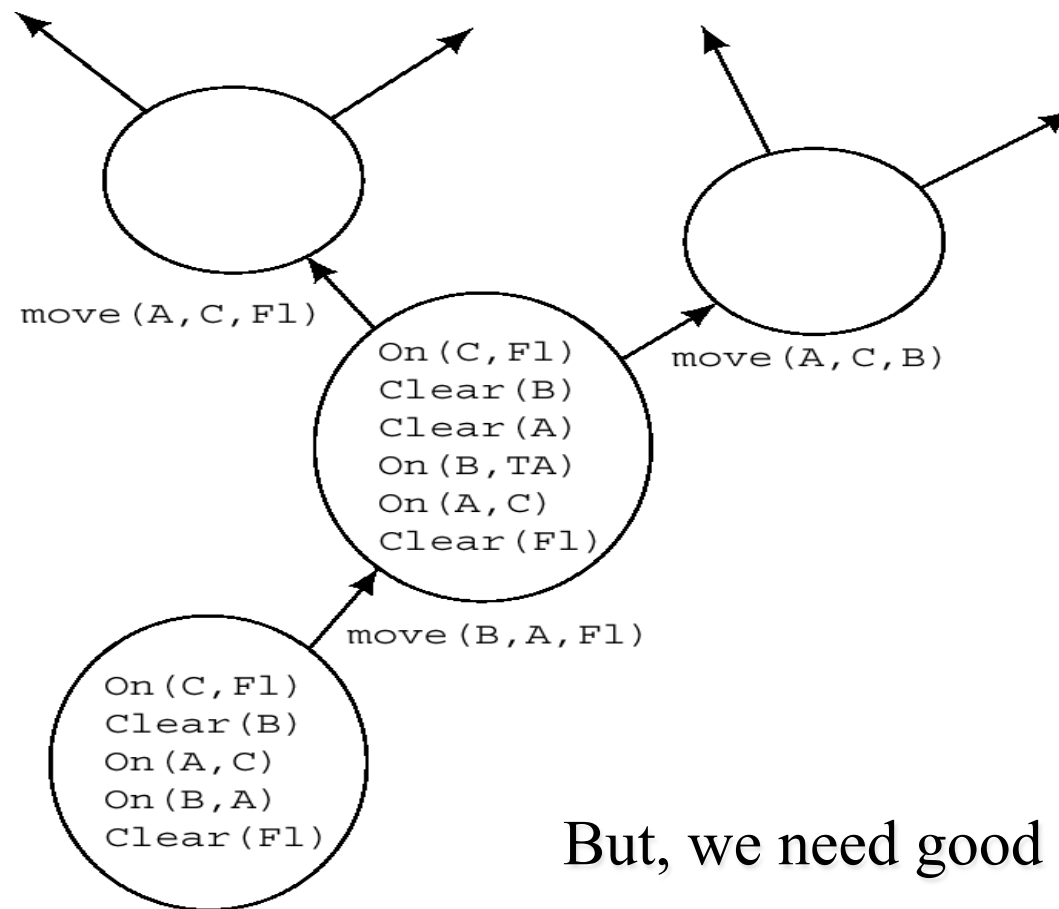
Algorithms for Planning as State-space

- Forward (progression) state-space
- Backward (regression) state-space search
- Heuristic search
- Planning graphs

Planning forward and backwards



Forward Search Methods:
can use A* with some h and g



But, we need good heuristics

Backward: Recursive STRIPS

- **Forward** search with islands:
- Achieve one subgoal at a time. Achieve a new conjunct without ever violating already achieved conjuncts or maybe temporarily violating previous subgoals.
- **General Problem Solver** (GPS) by Newell Shaw and Simon (1959) uses **Means-Ends** analysis.
- Each subgoal is achieved via a matched rule, then its preconditions are subgoals and so on. This leads to a planner called STRIPS(γ) when γ is a goal formula.

STRIPS algorithm


- Given a goal stack:
 1. Consider the top goal
 2. Find a sequence of actions satisfying the goal from the current state and apply them.
 3. The next goal is considered from the new state.
 4. Termination: stack empty
 5. Check goals again.

The Sussman anomaly

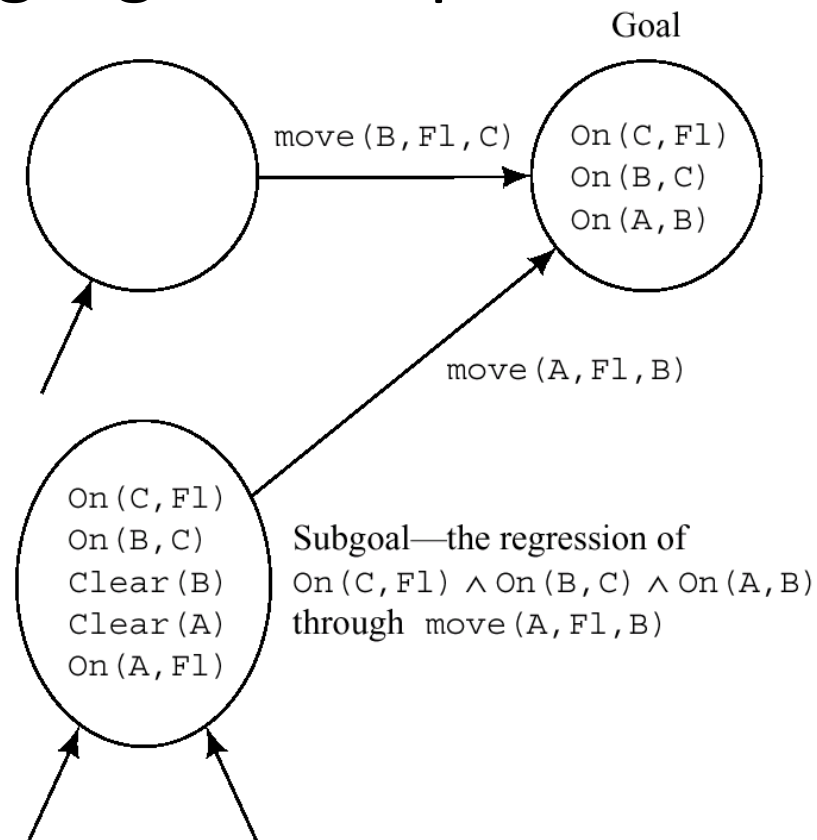
- RSTRIPS cannot achieve shortest plan
- Two possible orderings of subgoals:
 - On(A,B) and On(B,C) or On(B,C) and On(A,B)



Backward search methods;

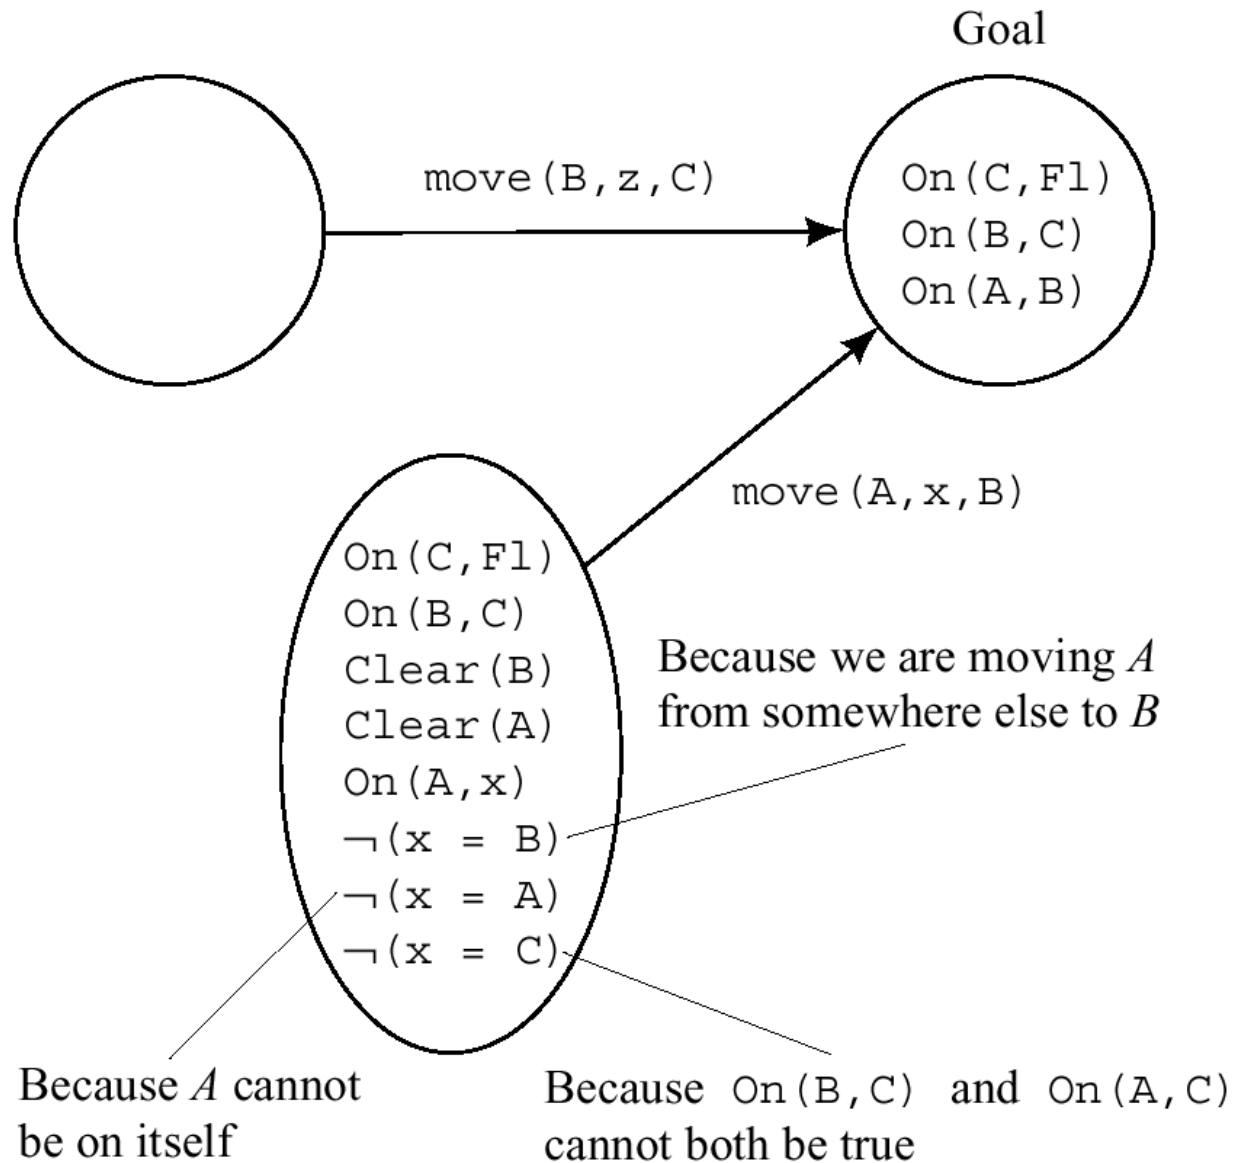


- Regressing a ground operator

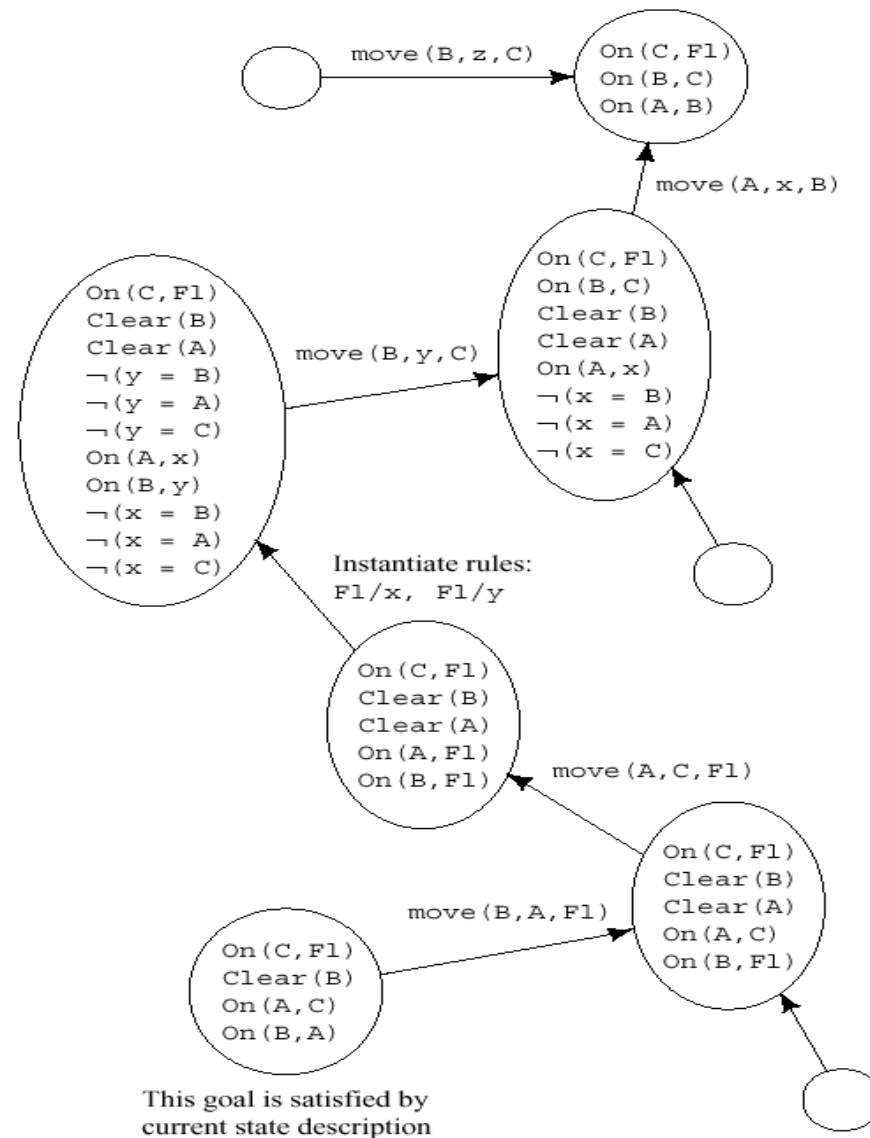


Continue until a subgoal is produced
that is satisfied by current world state

Regressing an ungrounded operator



Example of Backward Search



Heuristics for planning

- **Use relax problem idea** to get lower bounds on least number of actions to the goal.
 - Remove all or some preconditions
- **Sub-goal independence**: the cost of solving a set of subgoals equals the sum cost of solving each one independently, or max cost
 - Can be pessimistic (interacting sub-plans)
 - Can be optimistic (negative effects)
- **Simple**: number of unsatisfied sub-goals.
- Various ideas related to removing negative effects or positive effects.

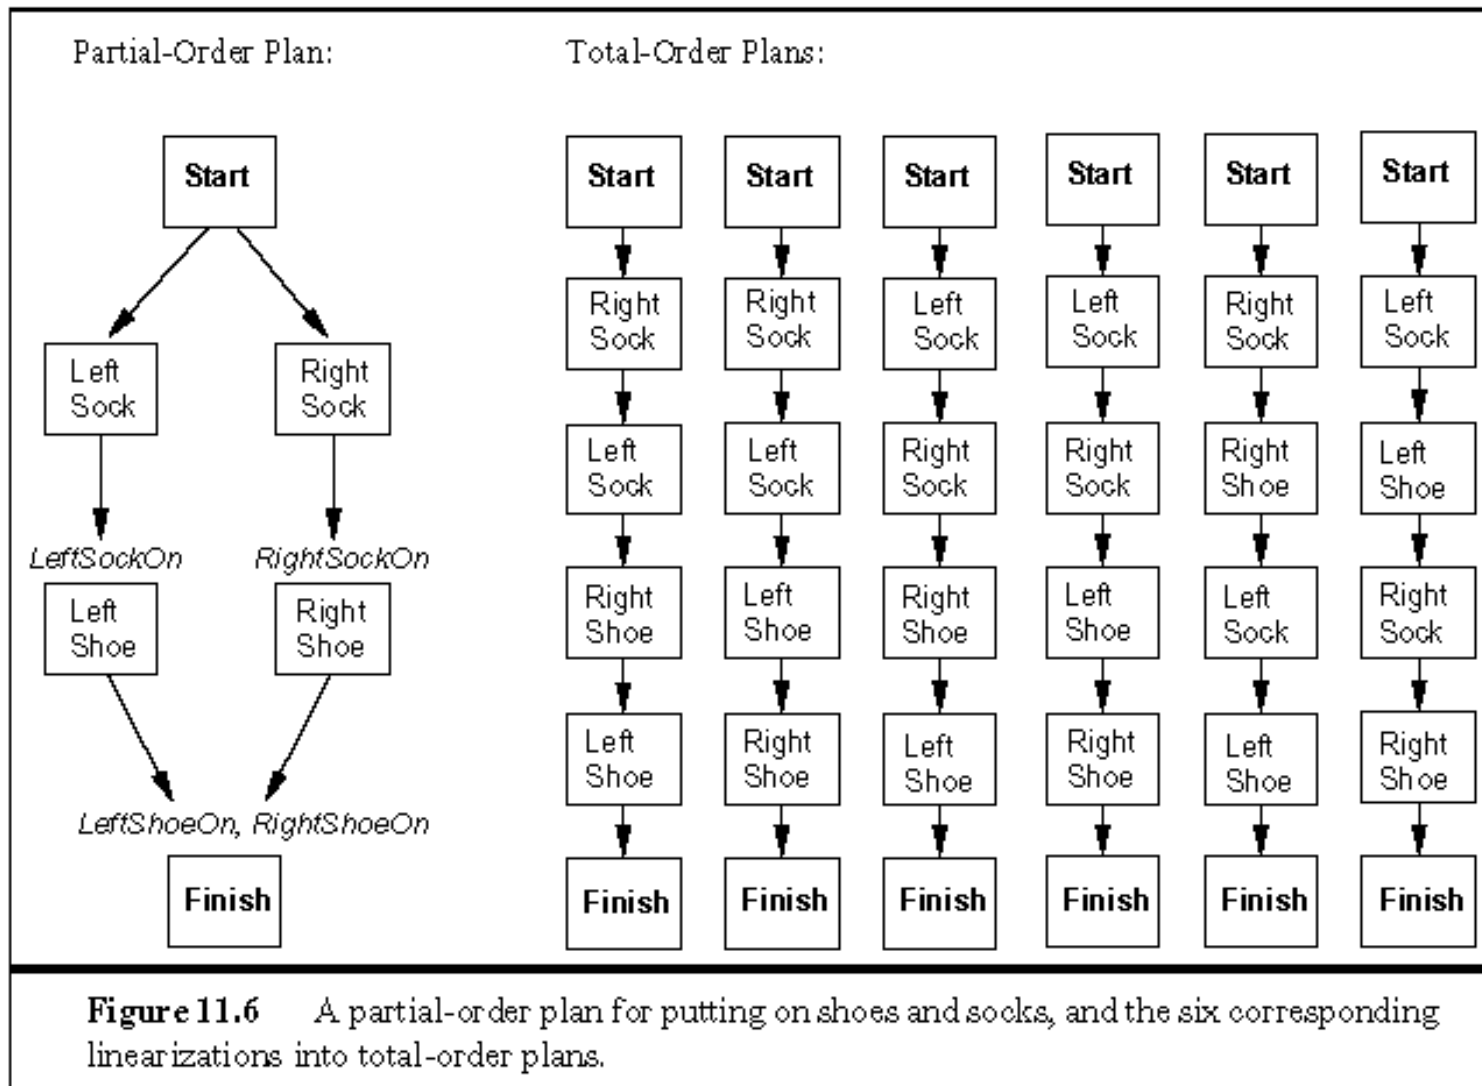
More on heuristic generation

- Ignore pre-conditions
- Ignore delete list: allow making monotone progress toward the goal.
 - Still NP-hard for optimal solution, but hill-climbing algorithms using an approximate solution that is polynomial. (example, 15 puzzle)
- Abstraction: Combines many states into a single one: Pattern databases
- FF : Fast-forward planner (Hoffman 2005), a forward state-space planner with delete-list based heuristic

Partial order planning

- Least commitment planning
- Nonlinear planning
- Search in the space of partial plans
- A state is a partial incomplete partially ordered plan
- Operators transform plans to other plans by:
 - Adding steps
 - Reordering
 - Grounding variables
- SNLP: Systematic Nonlinear Planning (McAllester and Rosenblitt 1991)
- NONLIN (Tate 1977)

A partial order plan for putting shoes and socks



Planning Graphs

- A planning graph consists of a sequence of levels that correspond to time-steps in the plan
- Level 0 is the initial state.
- Each level contains a set of literals and a set of actions
- Literals are those that could be true at the time step.
- Actions are those that their preconditions could be satisfied at the time step.
- Works only for propositional planning.

Example: Have cake and eat it too

```
Init{Have{Cake}}  
Goal{Have{Cake}  $\wedge$  Eaten{Cake}}  
Action{Eat{Cake}}  
  PRECOND: Have{Cake}  
  EFFECT:  $\neg$  Have{Cake}  $\wedge$  Eaten{Cake}}  
Action{Bake{Cake}}  
  PRECOND:  $\neg$  Have{Cake}  
  EFFECT: Have{Cake}
```

Figure 11.11 The “have cake and eat cake too” problem.

The Planning graphs for “have cake”,

- Persistence actions: Represent “inactions” by boxes: frame axiom
- Mutual exclusions (mutex) are represented between literals and actions.
- S_1 represents multiple states
- Continue until two levels are identical. The graph levels off.
- The graph records the impossibility of certain choices using mutex links.
- Complexity of graph generation: polynomial in number of literals.




Figure 11.12 The planning graph for the “have cake and eat cake too” problem up to level S_2 . Rectangles indicate actions (small squares indicate persistence actions) and straight lines indicate preconditions and effects. Mutex links are shown as curved gray lines.

Defining Mutex relations

- A mutex relation holds between two actions on the same level iff any of the following holds:
 - **Inconsistency effect:** one action negates the effect of another. Example “eat cake and persistence of have cake”
 - **Interference:** One of the effect of one action is the negation of the precondition of the other. Example: eat cake and persistence of Have cake
 - **Competing needs:** one of the preconditions of one action is mutually exclusive with a precondition of another. Example: Bake(cake) and Eat(Cake).
 - **A mutex relation holds between 2 literals** at the same level iff one is the negation of the other or if each possible pair of actions that can achieve the 2 literals is mutually exclusive.

Planning graphs for heuristic estimation

- Estimate the cost of achieving a goal by the level in the planning graph where it appears.
- To estimate the cost of a conjunction of goals use one of the following:
 - Max-level: take the maximum level of any goal (admissible)
 - Sum-cost: Take the sum of levels (inadmissible)
 - Set-level: find the level where they all appear without Mutex (admissible). Dominates max-level
- Graph plans are relaxation of the problem. Representing more than pair-wise mutex is not cost-effective

The graphplan algorithm

```
function GRAPHPLAN(problem) returns solution or failure

graph ← INITIAL-PLANNING-GRAPH(problem)
goals ← GOALS[problem]
loop do
  if goals all non-mutex in last level of graph then do
    solution ← EXTRACT-SOLUTION(graph, goals, LENGTH(graph))
    if solution ≠ failure then return solution
    else if NO-SOLUTION-POSSIBLE(graph) then return failure
  graph ← EXPAND-GRAPH(graph, problem)
```

Figure 11.13 The GRAPHPLAN algorithm. GRAPHPLAN alternates between a solution extraction step and a graph expansion step. EXTRACT-SOLUTION looks for whether a plan can be found, starting at the end and searching backwards. EXPAND-GRAPH adds the actions for the current level and the state literals for the next level.

Planning graph for spare tire

goal: $at(spare, axle)$

- S_2 has all goals and no mutex so we can try to extract solutions
- Use either CSP algorithm with actions as variables
- Or search backwards

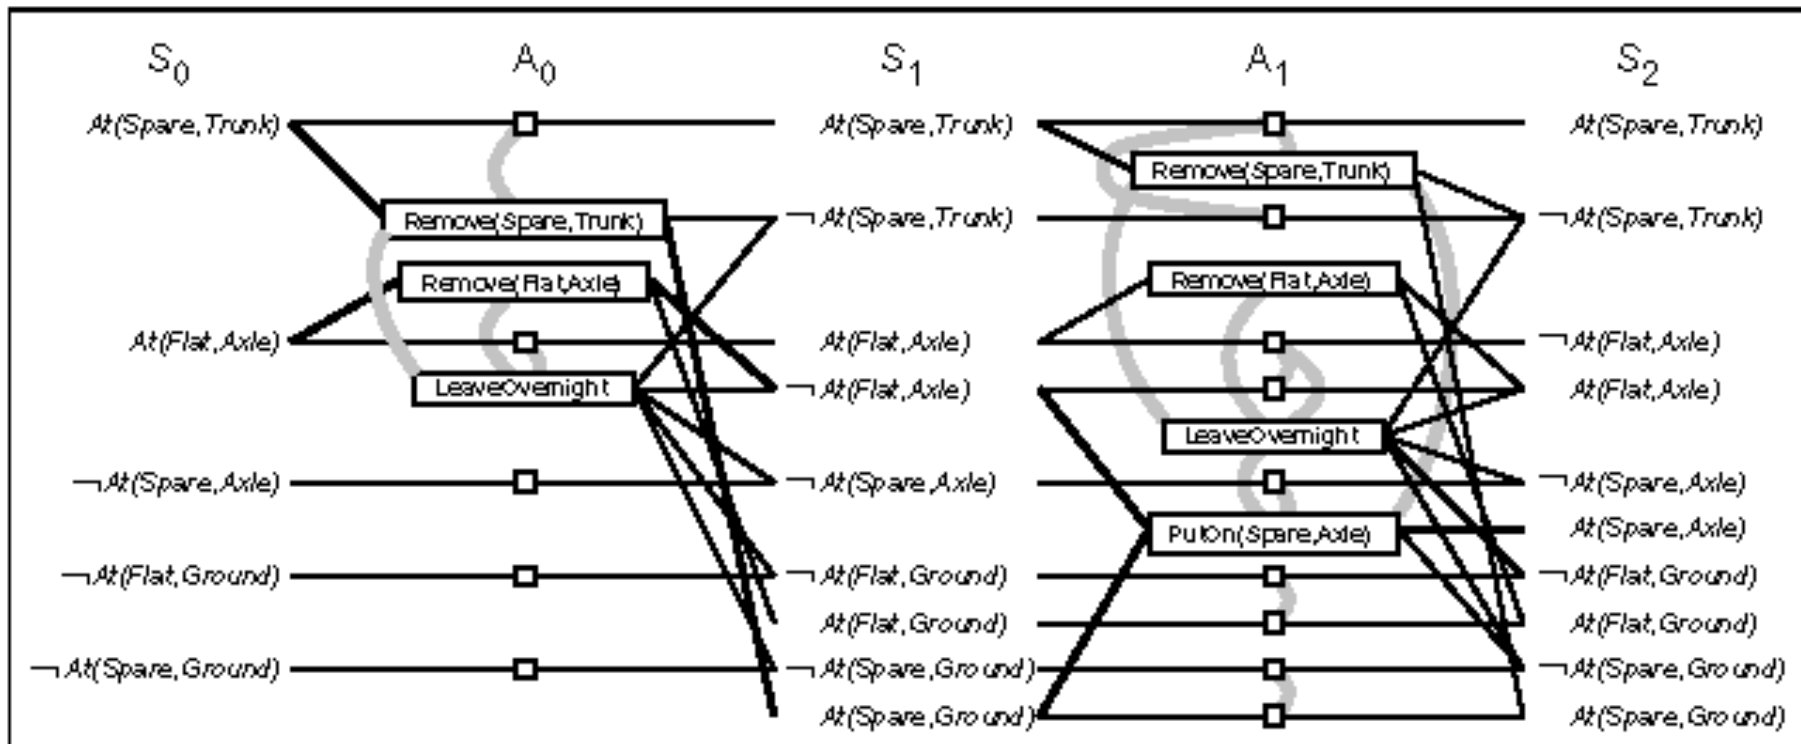


Figure 11.14 The planning graph for the spare tire problem after expansion to level S_2 . Mutex links are shown as gray lines. Only some representative mutexes are shown, because the graph would be too cluttered if we showed them all. The solution is indicated by bold lines and outlines.

Search planning-graph backwards with heuristics

- How to choose an action during backwards search:
 - Use greedy algorithm based on the level cost of the literals.
- For any set of goals:
 - 1. Pick first the literal with the highest level cost.
 - 2. To achieve the literal, choose the action with the easiest preconditions first (based on sum or max level of precond literals).

Properties of planning graphs; termination

- Literals increase monotonically
 - Once a literal is in a level it will persist to the next level
- Actions increase monotonically
 - Since the precondition of an action was satisfied at a level and literals persist the action's precond will be satisfied from now on
- Mutexes decrease monotonically:
 - If two actions are mutex at level S_i , they will be mutex at all previous levels at which they both appear
- Because literals increase and mutex decrease it is guaranteed that we will have a level where all goals are non-mutex

Other classical planning approaches

- The most effective approaches to planning currently are:
 - Translating to Boolean Satisfiability
 - Forward state-space search with carefully crafted heuristics
 - Search using planning graphs (covered already)

Planning as Satisfiability

- Express propositional planning as a set of propositions.
- Index propositions with time steps:
- $On(A,B)_0$, $ON(B,C)_0$
- Goal conditions: the goal conjuncts at time T , T is determined arbitrarily.
- Unknown propositions are not stated.
- Propositions known not to be true are stated negatively.
- Actions: a proposition for each action for each time slot.
- Successor state axioms need to be expressed for each action (like in the situation calculus but it is propositional)

Planning with propositional logic (continued)

- We write the formula:
 - Initial state and successor state axioms and goal
- We search for a model to the formula. Those actions that are assigned true constitute a plan.
- To have a single plan we may have a mutual exclusion for all actions in the same time slot.
- We can also choose to allow partial order plans and only write exclusions between actions that interfere with each other.
- Planning: iteratively try to find longer and longer plans.

SATplan algorithm

```
function SATPLAN(problem,  $T_{\text{max}}$ ) returns solution or failure
  inputs: problem, a planning problem
            $T_{\text{max}}$ , an upper limit for plan length

  for  $T = 0$  to  $T_{\text{max}}$  do
    cnf, mapping  $\leftarrow$  TRANSLATE-TO-SAT(problem,  $T$ )
    assignment  $\leftarrow$  SAT-SOLVER(cnf)
    if assignment is not null then
      return EXTRACT-SOLUTION(assignment, mapping)
  return failure
```

Figure 11.15 The SATPLAN algorithm. The planning problem is translated into a CNF sentence in which the goal is asserted to hold at a fixed time step T and axioms are included for each time step up to T . (Details of the translation are given in the text.) If the satisfiability algorithm finds a model, then a plan is extracted by looking at those proposition symbols that refer to actions and are assigned *true* in the model. If no model exists, then the process is repeated with the goal moved one step later.

Complexity of satplan

- The total number of action symbols is:
 - $|T| \times |Act| \times |O|^p$
 - O = number of objects, p is scope of atoms.
- Number of clauses is higher.
- Example: 10 time steps, 12 planes, 30 airports, the complete action exclusion axiom has 583 million clauses.

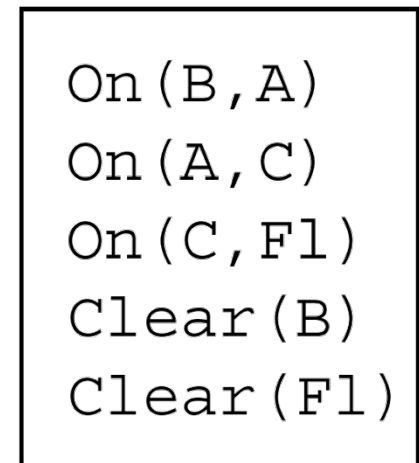

Summary: Planning

- STRIPS Planning
- Forward and backward planning
- Partial order planning
- Situation Calculus
- Heirarchical planning
- Graph planning
- Satplan
- Readings: RN chapter 10

The Situation Calculus

- A **goal** can be described by a sentence (wff) $On(x, B)$ if we want to have a block on B
 - **Planning**: finding a set of actions to achieve a goal wff.
 - **Situation Calculus** (McCarthy, Hayes, 1969, Green 1969)
 - A Predicate Calculus formalization of *states*, *actions*, and their *effects*.
- $On(B, A) \wedge On(A, C) \wedge On(C, F_1) \wedge Clear(B) \wedge clear(F_1)$
 – S_0 state in the figure can be described by:

we reify the state and include them as arguments



The Situation Calculus (continued)

- The atoms denotes relations over states called **fluents**.

$$O_n(B, A, S_0) \wedge O_n(A, C, S_0) \wedge O_n(C, F_1, S_0) \wedge clear(B, S_0)$$

- We can also have.

$$(\forall x, y, s)[On(x, y, s) \wedge \neg(y = F_1) \rightarrow \neg Clear(y, s)]$$

- Knowledge about state and actions \Rightarrow predicate calculus theory.
- Inference can be used to answer:
 - Is there a state satisfying a goal?
 - How can the present state be transformed into that state by actions? The answer is a *plan*

Representing Actions (continued)

- Express the effects of actions.
 - Example: (on, move) (expresses the effect of move on “On”)

$[On(x, y, s) \wedge Clear(x, s) \wedge Clear(z, s) \wedge (x \neq z)$

$\rightarrow On(x, z, do(move(x, y, z), s))]$

negative :

$[On(x, y, s) \wedge Clear(x, s) \wedge Clear(z, s) \wedge (x \neq z)$

$\rightarrow \neg On(x, y, do(move(x, y, z), s))]$

- Positive: describes how action makes a fluent true
- Negative : describes how action makes a fluent false
- Antecedent: pre-condition for actions
- Consequent: how the fluent is changed

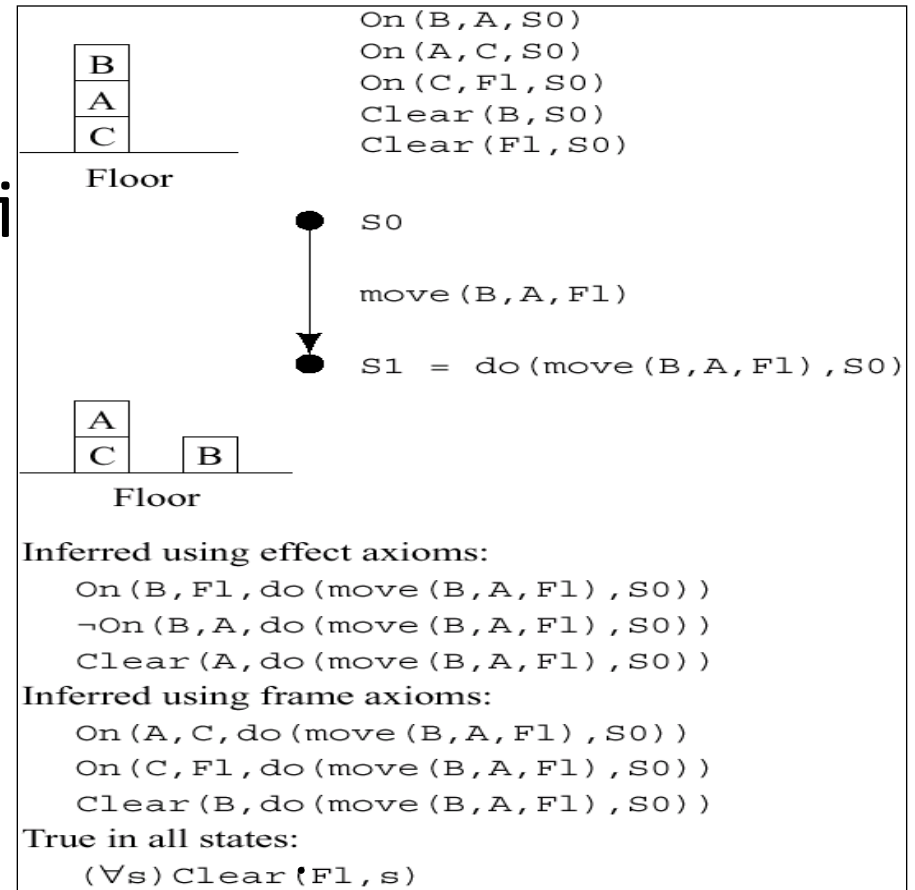
Representing Actions (continued)

- Effect axioms for (clear, move): (move(x,y,z))

- Precondition are satisfied

- B/x, A/y, S0/s, F1/z

- What was true in S0 remains true



Frame Axioms

- Not everything true can be inferred
On(C,F1) remains true but cannot be inferred
- Actions have local effect
 - We need **frame axioms** for each action and each fluent that does not change as a result of the action
 - example: frame axioms for (move, on)
 - If a block is on another block and *move* is not relevant, it will stay the same.

- Positive:

$$[On(x, y, s) \wedge (x \neq u)] \rightarrow On(x, y, do(move(u, v, z), s))$$

- negative

$$(\neg On(x, y, s) \wedge [(x \neq u) \vee (y \neq z)]) \rightarrow \neg On(x, y, do(move(u, v, z), s))$$

Frame Axioms (continued)

- Frame axioms for (move, clear):

$$Clear(u, s) \wedge (u \neq z) \rightarrow Clear(u, do(move(x, y, z), s))$$

$$\neg Clear(u, s) \wedge (u \neq y) \rightarrow \neg Clear(u, do(move(x, y, z), s))$$

- **The frame problem:** need axioms for every pair of {action, fluent}!!!
- There are languages that embed some assumption on frame axioms that can be derived automatically:
 - Default logic
 - Negation as failure
 - Nonmonotonic reasoning
 - Minimizing change

Other problems

- **The qualification problem:** qualifying the antecedent for all possible exception. Needs to enumerate all exceptions
 - \sim heavy and \sim glued and \sim armbroken \rightarrow can-move
 - \sim bird and \sim cast-in-concrete and \sim dead... \rightarrow flies
- Solutions: default logics, Non-monotonic logics
- **The ramification problem:**
 - If a robot carries a package, the package will be where the robot is. But what about the frame axiom, when can we infer about the effect of the actions and when we cannot.

Generating plans

- To generate a plan to achieve a goal, we attempt to prove γ

$$\exists s \gamma(s)$$

- Example: Get block B on the floor from S0. $\exists s \text{On}(B, F_1, s)$
- Prove:
- By resolution refutation: add $\exists s \neg \text{On}(B, F_1, s)$ for all s not $\text{On}(B, F_1, s)$
- (page 370 top, try it)

Proof

- $\sim \text{On}(B, F1, s) \vee \text{Ans}(s)$
- $\text{On}(B, A, S0)$
- $\text{On}(A, C, S0)$
- $\text{On}(C, F1, S0)$
- $\text{Clear}(B, S0)$
- $\text{Clear}(F1, S0)$
- $[\text{On}(x, y, s) \wedge \text{Clear}(x, s) \wedge \text{Clear}(z, s) \wedge (x \text{ not equal } y) \rightarrow \text{On}(x, z, \text{do}(\text{move}(x, y, z), s))]$
- (inequality is negation of equality and we assume infinite supply of $(A=A)$)