

State Spaces & Partial-Order Planning

AI Class 22 (Ch. 10 through 10.4.4)

Material from Dr. Marie desJardins. Some material adapted from notes by Andreas Geyer-Schulz and Chuck Dyer

Overview

- What is planning?
- Approaches to planning
 - GPS / STRIPS
 - Situation calculus formalism [revisited]
 - Partial-order planning

Planning Problem

- What is the planning problem?
- Find a **sequence of actions** that achieves a **goal** when executed from an **initial state**.
- That is, given
 - A set of operators (possible actions)
 - An initial state description
 - A goal (description or conjunction of predicates)
- Compute a sequence of operations: a **plan**.

Planning Problem

- What is the planning problem?
 - Find a **sequence of actions** that achieves a **goal** when executed from an **initial state**.
 - That is, given
 - A set of operators (possible actions)
 - An initial state description
 - A goal (description or conjunction of predicates)
 - Compute a sequence of operations:
-
- pants off
 - right shoe off
 - right sock off
 - right shoe off (etc)
 - pants on (etc)

Typical Assumptions (1)

- **Atomic time**: Each action is indivisible
 - Can't be interrupted halfway through putting on pants
- **No concurrent actions** allowed
 - Can't put on socks at the same time
- **Deterministic actions**
 - The result of actions are completely known – no uncertainty

Typical Assumptions

- Agent is the **sole cause of change** in the world
 - Nobody else is putting on your socks
- Agent is **omniscient**:
 - Has complete knowledge of the state of the world
- **Closed world assumption**:
 - Everything known-true about the world is in the *state description*
 - Anything not known-true is known-false

Blocks World


The **blocks world** consists of a table, set of blocks, and a robot gripper

Some domain constraints:

- Only one block on another block
- Any number of blocks on table
- Hand can only hold one block


Typical representation:

`ontable(a)` `handempty`
`ontable(c)` `on(b,a)`
`clear(b)` `clear(c)`



Blocks world

- A micro-world
- Some domain constraints:
 - Only one block can be on another block
 - Any number of blocks can be on the table
 - The hand can only hold one block



Meant to be a simple model!
 Try demo at:
<http://aispace.org/planning/>


Typical BW planning problem

Initial state:

`clear(a)`
`clear(b)`
`clear(c)`
`ontable(a)`
`ontable(b)`
`ontable(c)`
`handempty`

Goal state:

`on(b,c)`
`on(a,b)`
`ontable(c)`




Typical BW planning problem

Initial state:

`clear(a)`
`clear(b)`
`clear(c)`
`ontable(a)`
`ontable(b)`
`ontable(c)`
`handempty`

Goal state:

`on(b,c)`
`on(a,b)`
`ontable(c)`



assertions describing a state

atomic robot actions

Plan:
`pickup(b)`
`stack(b,c)`
`pickup(a)`
`stack(a,b)`

Major Approaches

- GPS / STRIPS
- **Situation calculus**
- **Partial order planning**
- Hierarchical decomposition (HTN planning)
- Planning with constraints (SATplan, Graphplan)
- **Reactive planning**

Planning vs. problem solving

- Planning vs. problem solving: can often solve similar problems
- Planning is more powerful and efficient because of the representations and methods used
- States, goals, and actions are decomposed into sets of sentences (usually in first-order logic)
- Search often proceeds through *plan space* rather than *state space* (though there are also state-space planners)
- Sub-goals can be planned independently, reducing the complexity of the planning problem

Another BW planning problem

Initial state:

- clear(a)
- clear(b)
- clear(c)
- ontable(a)
- ontable(b)
- ontable(c)
- handempty

Goal:

- on(a,b)
- on(b,c)
- ontable(c)

A plan

- pickup(a)
- stack(a,b)
- unstack(a,b)
- putdown(a)
- pickup(b)
- stack(b,c)
- pickup(a)
- stack(a,b)

Yet Another BW planning problem

Initial state:

- clear(c)
- ontable(a)
- on(b,a)
- on(c,b)
- handempty

Goal:

- on(a,b)
- on(b,c)
- ontable(c)

Plan:

- unstack(c,b)
- putdown(c)
- unstack(b,a)
- putdown(b)
- putdown(b)
- pickup(a)
- stack(a,b)
- unstack(a,b)
- putdown(a)
- pickup(b)
- stack(b,c)
- pickup(a)
- stack(a,b)

backtracking

Planning as Search

- Can think of planning as a search problem
 - **Actions:** generate successor states
 - **States:** completely described & only used for successor generation, heuristic fn. evaluation & goal testing
 - **Goals:** represented as a goal test and using a heuristic function
 - **Plan representation:** unbroken sequences of actions forward from initial states or backward from goal state

“Get a quart of milk, a bunch of bananas and a variable-speed cordless drill.”

Slightly more complex KB:

Treating planning as a search problem isn't very efficient!

General Problem Solver

- The **General Problem Solver (GPS)** system
 - An early planner (Newell, Shaw, and Simon)
- Generate actions that *reduce difference* between current state and goal state
- Uses *Means-Ends Analysis*
 - Compare what is **given** or **known** with what is desired
 - Select a reasonable thing to do next
 - Use a **table of differences** to identify procedures to reduce differences
- GPS is a **state space planner**
 - Operates on state space problems specified by an initial state, some goal states, and a set of operations

Situation Calculus Planning

- Intuition: Represent the **planning problem** using first-order logic
 - Situation calculus lets us reason about **changes** in the world
 - Use theorem proving to show (“prove”) that a sequence of actions will lead to a desired result, when applied to a world state / situation

Situation Calculus Planning, cont.

- **Initial state:** a logical sentence about (situation) S_0
- **Goal state:** usually a conjunction of logical sentences
- **Operators:** descriptions of how the world changes as a result of the agent's actions:
 - $\text{Result}(a,s)$ names the situation resulting from executing action a in situation s .
- Action sequences are also useful:
 - $\text{Result}(l,s)$: result of executing list of actions l starting in s

Situation Calculus Planning, cont.

- **Initial state:**
 $\text{At}(\text{Home}, S_0) \wedge \neg \text{Have}(\text{Milk}, S_0) \wedge \neg \text{Have}(\text{Bananas}, S_0) \wedge \neg \text{Have}(\text{Drill}, S_0)$
- **Goal state:**
 $(\exists s) \text{At}(\text{Home}, s) \wedge \text{Have}(\text{Milk}, s) \wedge \text{Have}(\text{Bananas}, s) \wedge \text{Have}(\text{Drill}, s)$
- **Operators:**
 $\forall (a,s) \text{Have}(\text{Milk}, \text{Result}(a,s)) \Leftrightarrow ((a = \text{Buy}(\text{Milk}) \wedge \text{At}(\text{Grocery}, s)) \vee (\text{Have}(\text{Milk}, s) \wedge a \neq \text{Drop}(\text{Milk})))$
- **Result(a,s):** situation after executing action a in situation s
 $(\forall s) \text{Result}([], s) = s$
 $(\forall a,p,s) \text{Result}([a|p]s) = \text{Result}(p, \text{Result}(a,s))$

p=plan

Situation Calculus, cont.

- Solution: a **plan** that when applied to the **initial state** gives a situation satisfying the **goal query**:
 $\text{At}(\text{Home}, \text{Result}(p, S_0))$
 $\wedge \text{Have}(\text{Milk}, \text{Result}(p, S_0))$
 $\wedge \text{Have}(\text{Bananas}, \text{Result}(p, S_0))$
 $\wedge \text{Have}(\text{Drill}, \text{Result}(p, S_0))$
- Thus we would expect a plan (i.e., variable assignment through unification) such as:
 $p = [\text{Go}(\text{Grocery}), \text{Buy}(\text{Milk}), \text{Buy}(\text{Bananas}), \text{Go}(\text{HardwareStore}), \text{Buy}(\text{Drill}), \text{Go}(\text{Home})]$

Situation Calculus: Blocks World

- Example situation calculus rule for blocks world:
 $\text{clear}(X, \text{Result}(A,S)) \Leftrightarrow$
 $[\text{clear}(X, S) \wedge$
 $(\neg (A = \text{Stack}(Y,X) \vee A = \text{Pickup}(X))$
 $\vee (A = \text{Stack}(Y,X) \wedge \neg (\text{holding}(Y,S)))$
 $\vee (A = \text{Pickup}(X) \wedge \neg (\text{handempty}(S) \wedge \text{ontable}(X,S) \wedge \text{clear}(X,S)))]]$
 $\vee [A = \text{Stack}(X,Y) \wedge \text{holding}(X,S) \wedge \text{clear}(Y,S)]$
 $\vee [A = \text{Unstack}(Y,X) \wedge \text{on}(Y,X,S) \wedge \text{clear}(Y,S) \wedge \text{handempty}(S)]$
 $\vee [A = \text{Putdown}(X) \wedge \text{holding}(X,S)]]$
- English translation: a block is **clear** if
 - in the previous state it was clear AND we didn't pick it up or stack something on it successfully, or
 - we stacked it on something else successfully, or
 - something was on it that we unstacked successfully, or
 - we were holding it and we put it down.

Wow.

Situation Calculus Planning: Analysis

- Fine in theory, but:
 - Problem solving (search) is exponential in the worst case
 - Resolution theorem proving only finds a proof (plan), not necessarily a *good* plan
- So what can we do?
 - Restrict the language
 - Blocks world is already pretty small...
 - Use a **special-purpose planner rather than general theorem prover**

Basic Representations for Planning

- Classic approach first used in the STRIPS planner circa 1970
- **States** represented as conjunction of ground literals
 $\text{at}(\text{Home}) \wedge \neg \text{have}(\text{Milk}) \wedge \neg \text{have}(\text{bananas}) \dots$
- Goals are conjunctions of literals, but may have variables*
 $\text{at}(?x) \wedge \text{have}(\text{Milk}) \wedge \text{have}(\text{bananas}) \dots$
- Don't need to fully specify state
 - Un-specified: either don't-care or assumed-false
 - Represent many cases in small storage
 - Often only represent **changes in state** rather than entire situation
- Unlike theorem prover, not finding whether the goal is **true**, but whether there is a sequence of actions to attain it

*generally assume \exists

Operator/Action Representation

- **Operators** contain three components:
 - **Action description**
 - **Precondition** - conjunction of positive literals
 - **Effect** - conjunction of positive or negative literals which describe how situation changes when operator is applied
- Example:

Op(Action: Go(there), Path(there,there), At(there), ~At(there))

Precond: At(there) \wedge Path(there,there), Effect: At(there) \wedge \sim At(there)

Go(there)
- All variables are **universally** quantified
- Situation variables are implicit
 - **Preconditions** must be true in the state immediately before operator is applied
 - **Effects** are true immediately after

Blocks World Operators

- Classic basic **operations** for the blocks world:
 - **stack(X,Y)**: put block X on block Y
 - **unstack(X,Y)**: remove block X from block Y (we saw these implicitly in the examples)
 - **pickup(X)**: pickup block X
 - **putdown(X)**: put block X on the table
- Each will be represented by
 - Preconditions
 - New facts to be added (add-effects)
 - Facts to be removed (delete-effects)
 - A set of (simple) variable constraints (optional!)

Blocks World Operators

- So given these operations:
 - **stack(X,Y), unstack(X,Y), pickup(X), putdown(X)**
- Need:
 - **Preconditions**, facts to be added (**add-effects**), facts to be removed (**delete-effects**), optional variable constraints

Example: stack

preconditions(stack(X,Y), [holding(X), clear(Y)])

deletes(stack(X,Y), [holding(X), clear(Y)]).

adds(stack(X,Y), [handempty, on(X,Y), clear(X)])

constraints(stack(X,Y), [X \neq Y, Y \neq table, X \neq table])

Blocks World Operators II

operator(stack (X,Y), Precond [holding(X), clear(Y)], Add [handempty, on(X,Y), clear(X)], Delete [holding(X), clear(Y)], Constr [X \neq Y, Y \neq table, X \neq table]).	operator(unstack (X,Y), [on(X,Y), clear(X), handempty], [holding(X), clear(Y)], [handempty, clear(X), on(X,Y)], [X \neq Y, Y \neq table, X \neq table]).
operator(pickup (X), [ontable(X), clear(X), handempty], [holding(X)], [ontable(X), clear(X), handempty], [X \neq table]).	operator(putdown (X), [holding(X)], [ontable(X), handempty, clear(X)], [holding(X)], [X \neq table]).


Plan-Space Planning

- Alternative: **search through space of plans**, not situations
- Start from a **partial plan**; expand and refine until a complete plan that solves the problem is generated
- **Refinement operators** add constraints to the partial plan and modification operators for other changes
- We can still use STRIPS-style operators:
 - Op(ACTION: PutOnRightShoe, PRECOND: RightSockOn, EFFECT: RightShoeOn)
 - Op(ACTION: PutOnRightSock, EFFECT: RightSockOn)
 - Op(ACTION: PutOnLeftShoe, PRECOND: LeftSockOn, EFFECT: LeftShoeOn)
 - Op(ACTION: PutOnLeftSock, EFFECT: LeftSockOn)

Partial-Order Planning


Partial-Order Planning

- The big idea: Don't specify the order of steps if you don't have to.




- Doesn't matter, but a regular planner has to consider and specify all the options.

A simple graphical notation



Partial-Order Planning

- A **linear planner** builds a plan as a **totally ordered sequence** of plan steps
- A **non-linear planner (aka partial-order planner)** builds up a plan as a set of steps with some temporal constraints
 - E.g., S1 < S2 (step S1 must come before S2)

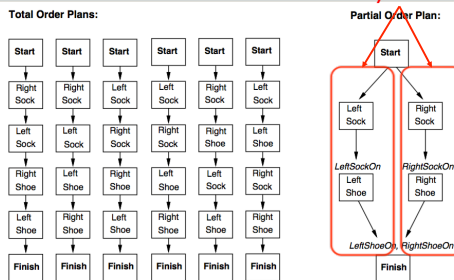


The order here does matter, so the planner has to know that.

- Partially ordered plan (POP) **refined** by either:
 - adding a new **plan step**, or
 - adding a new **constraint** to the steps already in the plan.
- A POP can be linearized by topological sorting – R&N 223

Linear vs. POP: Shoes

Do these sequences in any order




Some example domains

- We'll use some simple problems to illustrate planning problems and algorithms
- Putting on your socks and shoes in the morning
 - Actions like put-on-left-sock, put-on-right-shoe
- Planning a shopping trip involving buying several kinds of items
 - Actions like go(X), buy(Y)

The Initial Plan

Every plan starts the same way



Least Commitment

- Non-linear planners embody the principle of **least commitment**
 - Only choose actions, orderings and variable bindings absolutely necessary, postponing other decisions
 - Avoid early commitment to decisions that don't really matter
- Linear planners always choose to add a plan step in a particular place in the sequence
- Non-linear planners choose to add a step and possibly some temporal constraints

Non-Linear Plan Components

- 1) A set of **steps** $\{S_1, S_2, S_3, S_4, \dots\}$
 - Each step has an **operator description, preconditions and post-conditions**
 - **ACTION:** LeftShoe, **PRECOND:** LeftSockOn, **EFFECT:** LeftShoeOn
- 2) A set of **causal links** $\{ \dots (S_i, C, S_j) \dots \}$
 - (One) goal of step S_i is to achieve precondition C of step S_j
 - **(PutOnLeftShoe, LeftShoeOn, Finish)**
 - This says: No action that undoes LeftShoeOn is allowed to happen after PutOnLeftShoe and before Finish. Any action that undoes LeftShoeOn must either be before PutOnLeftShoe or after Finish.
- 3) A set of **ordering constraints** $\{ \dots S_i < S_j \dots \}$
 - If step S_i must come before step S_j
 - PutOnSock < Finish

Non-Linear Plan: Completeness

- A non-linear plan consists of
 - (1) A set of **steps** $\{S_1, S_2, S_3, S_4, \dots\}$
 - (2) A set of **causal links** $\{ \dots (S_i, C, S_j) \dots \}$
 - (3) A set of **ordering constraints** $\{ \dots S_i < S_j \dots \}$
- A non-linear plan is **complete** iff
 - Every step mentioned in (2) and (3) is in (1)
 - If S_j has prerequisite C , then there exists a causal link in (2) of the form (S_i, C, S_j) for some S_i
 - If (S_i, C, S_j) is in (2) and step S_k is in (1), and S_k threatens (S_i, C, S_j) (makes C false), then (3) contains either $S_k < S_i$ or $S_j < S_k$

Trivial Example

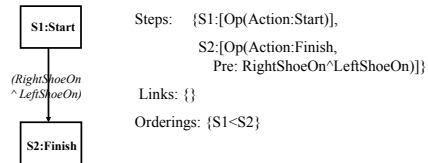
Operators:

Op(ACTION: RightShoe, PRECOND: RightSockOn, EFFECT: RightShoeOn)


Op(ACTION: RightSock, EFFECT: RightSockOn)

Op(ACTION: LeftShoe, PRECOND: LeftSockOn, EFFECT: LeftShoeOn)

Op(ACTION: LeftSock, EFFECT: leftSockOn)



Solution





POP Constraints and Search Heuristics


- Only add steps that reach a not-yet-achieved precondition
- Use a least-commitment approach:
 - Don't order steps unless they need to be ordered
- Honor causal links $S_1 \Rightarrow S_2$ that **protect** a condition c :
 - Never add an intervening step S_3 that violates c
 - If a parallel action **threatens** c (i.e., has the effect of negating or **cllobbering** c), resolve that threat by adding ordering links:
 - Order S_3 before S_1 (**demotion**)
 - Order S_3 after S_2 (**promotion**)

Partial-Order Planning Example

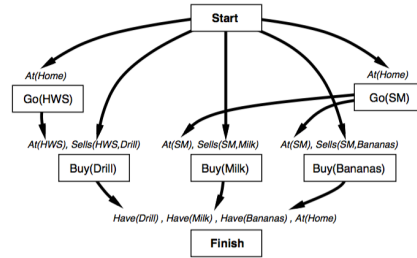

- **Initially:** at home; SM sells bananas; SM sells milk; HWS sells drills
- **Goal:** Be home with milk, bananas, and a drill



- Add three actions to achieve basic goals
- Use initial state to achieve the "Sells" preconditions
- Bold links are causal (protected), regular are just ordering constraints




Planning



Resolving Threats

- The S_3 action **threatens** the c precondition of S_2 if S_3 neither precedes nor follows S_2 and S_3 negates c .
 - We don't want to go to the HWS then leave before buying a drill...



Real-World Planning Domains

- Real-world domains are complex
 - Don't satisfy assumptions of STRIPS or partial-order planning methods
 - Some of the characteristics we may need to deal with:
 - Modeling and reasoning about resources
 - Representing and reasoning about time
 - Planning at different levels of abstractions
 - Conditional outcomes of actions
 - Uncertain outcomes of actions
 - Exogenous events
 - Incremental plan development
 - Dynamic real-time replanning
- } Scheduling
} Planning under uncertainty
} HTN planning

Hierarchical Planning

Hierarchical Decomposition

- The big idea: **Plan over high-level actions (HLAs), then figure out the steps to accomplish those.**
- Reduces complexity of planning space
 - Consider plan made of HLAs
 - **Then** make a plan for steps within each
 - Don't consider silly orderings that violate high-level concepts
- Can nest more than one level

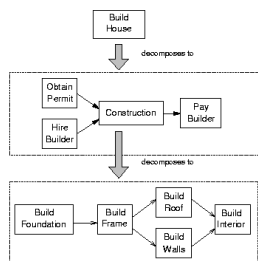
Hierarchical Decomposition: Example

- If we want to go to Hawaii (and we do)
 - Operators, unordered (because we haven't planned yet): [DriveToAirport](#), [TaxiToHotel](#), [PutClothesInSuitcase](#), [BuySungreen](#), [BoardPlane](#), [BuySwimsuit](#), [FindPassport](#), [PutPassportInCarryon](#), [DisembarkFromPlane](#), [BookHotel](#), ...
- High-Level Actions (HLAs): "Get to island" "Prepare for trip"
 - Order HLAs first: [PrepareForTrip](#) → [GetToIsland](#)
 - THEN order the subgoals within them
 - Don't have to consider "disembark" ↔ "find passport" ordering
- Nest as as needed
 - [PrepareForTrip](#) can include [ShopForTrip](#), which includes ...

Hierarchical Decomposition

- Hierarchical decomposition, or hierarchical task network (HTN) planning, uses **abstract operators** to **incrementally** decompose a planning problem from a **high-level goal** statement to a **primitive plan network**
- **Primitive operators** represent actions that are **executable**, and can appear in the final plan
- **Non-primitive operators** represent **goals** (equivalently, **abstract actions**) that require further decomposition (or *operationalization*) to be executed
- There is no "right" set of primitive actions: One agent's goals are another agent's actions!

HTN Planning: Example



HTN Operator: Example

```

OPERATOR decompose
PURPOSE: Construction
CONSTRAINTS:
  Length (Frame) <= Length (Foundation),
  Strength (Foundation) > Wt (Frame) + Wt (Roof)
  + Wt (Walls) + Wt (Interior) + Wt (Contents)
PLOT: Build (Foundation)
      Build (Frame)
      PARALLEL
        Build (Roof)
        Build (Walls)
      END PARALLEL
      Build (Interior)
  
```

HTN Operator Representation

- Russell & Norvig explicitly represent causal links
 - Can also be computed dynamically by using a model of preconditions and effects
 - Dynamically computing causal links means that actions from one operator can safely be interleaved with other operators, and subactions can safely be removed or replaced during plan repair
- R&N representation only includes variable bindings
 - Can actually introduce a wide array of variable constraints

Truth Criterion

- Determining whether a **formula is true** at a particular point in a partially ordered plan is, in the general case, NP-hard
- Intuition: there are exponentially many ways to **linearize** a partially ordered plan
- In the worst case, if there are N actions unordered with respect to each other, there are N! linearizations
- Ensuring soundness of truth criterion requires checking the formula under all possible linearizations
- Use heuristic methods instead to make planning feasible
- Check later to be sure no constraints have been violated

Truth Criterion in HTN Planners

- Heuristic:
 1. Prove that there exists *one* possible ordering of the actions that makes the formula true
 2. But don't insert ordering links to enforce that order
- Such a proof is efficient
 - Suppose you have an action A1 with a precondition P
 - Find an action A2 that achieves P (A2 can be initial world state)
 - Make sure there is no action *necessarily* between A2 and A1 that negates P
- Applying this heuristic for all preconditions in the plan can result in infeasible plans

Increasing Expressivity

- Conditional effects
 - Instead of different operators for different conditions, use a single operator with conditional effects
 - Move (block1, from, to) and MoveToTable (block1, from) collapse into one Move (block1, from, to):
 - Op(ACTION: Move(block1, from, to),
PRECOND: On (block1, from) ^ Clear (block1) ^ Clear (to)
EFFECT: On (block1, to) ^ Clear (from) ^ ~On(block1, from) ^ ~Clear(to) when to<>Table
 - There's a problem with this operator: can you spot it?
- Negated and disjunctive goals
- Universally quantified preconditions and effects

Reasoning About Resources

- What if I only have so much money for bananas and drills?
 - It suddenly matters that I don't introduce, e.g., [BuyGrapes](#)
- Introduce numeric variables that can be used as *measures*
- These variables represent resource quantities, and change over the course of the plan
- Certain actions **produce** (increase the quantity of) resources
- Other actions **consume** (decrease the quantity of) resources
- More generally, may want different types of resources
 - Continuous vs. discrete
 - Sharable vs. nonsharable
 - Reusable vs. consumable vs. self-replenishing

Other Real-World Planning Issues

- Conditional planning
- Partial observability
- Information gathering actions
- Execution monitoring and replanning
- Continuous planning
- Multi-agent (cooperative or adversarial) planning

POP Summary

- **Advantages**
 - Partial order planning is **sound** and **complete**
 - Typically produces **optimal** solutions (plan length)
 - Least commitment may lead to shorter search times
- **Disadvantages**
 - Significantly more complex algorithms
 - Hard to determine what is true in a state
 - Larger search space, since concurrent actions are allowed

Planning Summary

- **Planning representations**
 - Situation calculus
 - STRIPS representation: Preconditions and effects
- **Planning approaches**
 - State-space search (STRIPS, forward chaining, ...)
 - Plan-space search (partial-order planning, HTNs, ...)
 - *Constraint-based search (GraphPlan, SATplan, ...)*
- **Search strategies**
 - Forward planning
 - Goal regression
 - Backward planning
 - Least-commitment
 - Nonlinear planning