


Planning (Chapter 10)



Planning


- **Example problem:** *I'm at home and I need milk, bananas, and a drill. What do I do?*
- How is planning different from regular search?
 - States and action sequences typically have complex internal structure
 - State space and branching factor are huge
 - Multiple subgoals at multiple levels of resolution
- Examples of planning applications
 - Scheduling of tasks in space missions
 - Logistics planning for the army
 - Assembly lines, industrial processes
 - Robotics
 - Games, storytelling

A representation for planning


- STRIPS (Stanford Research Institute Problem Solver): classical planning framework from the 1970s
- **States** are specified as conjunctions of predicates
 - Start state: $At(Home) \wedge Sells(SM, Milk) \wedge Sells(SM, Bananas) \wedge Sells(HW, Drill)$
 - Goal state: $At(Home) \wedge Have(Milk) \wedge Have(Banana) \wedge Have(Drill)$
- **Actions** are described in terms of preconditions and effects:
 - $Go(x, y)$
 - **Precond:** $At(x)$
 - **Effect:** $\neg At(x) \wedge At(y)$
 - $Buy(x, store)$
 - **Precond:** $At(store) \wedge Sells(store, x)$
 - **Effect:** $Have(x)$
- Planning is “just” a search problem

Challenges of planning: “Sussman anomaly”


Start state:




Let's try to achieve
 $\text{On}(A,B)$:



Goal state:



Let's try to achieve
 $\text{On}(B,C)$:




Challenges of planning: “Sussman anomaly”

- Shows the limitations of *non-interleaved* planners that consider subgoals in sequence and try to satisfy them one at a time
 - If you try to satisfy subgoal X and then subgoal Y, X might undo some preconditions for Y, or Y might undo some effects of X
- More powerful planning approaches must *interleave* the steps towards achieving multiple subgoals

Algorithms for planning

- **Forward (progression) state-space search:**
starting with the start state, find all applicable actions
(actions for which preconditions are satisfied),
compute the successor state based on the effects,
keep searching until goals are met
 - Can work well with good heuristics



Algorithms for planning

- **Forward (progression) state-space search:**
starting with the start state, find all applicable actions (actions for which preconditions are satisfied), compute the successor state based on the effects, keep searching until goals are met
 - Can work well with good heuristics
- **Backward (regression) relevant-states search:**
to achieve a goal, what must have been true in the previous state?


Situation space planning vs. plan space planning

- **Situation space planners:** each node in the search space represents a world state, arcs are actions in the world
 - Plans are sequences of actions from start to finish
 - Must be *totally ordered*
- **Plan space planners:** nodes are (incomplete) plans, arcs are transformations between plans
 - Actions in the plan may be *partially ordered*
 - *Principle of least commitment:* avoid ordering plan steps unless absolutely necessary


Partial order planning

- Task: put on socks and shoes

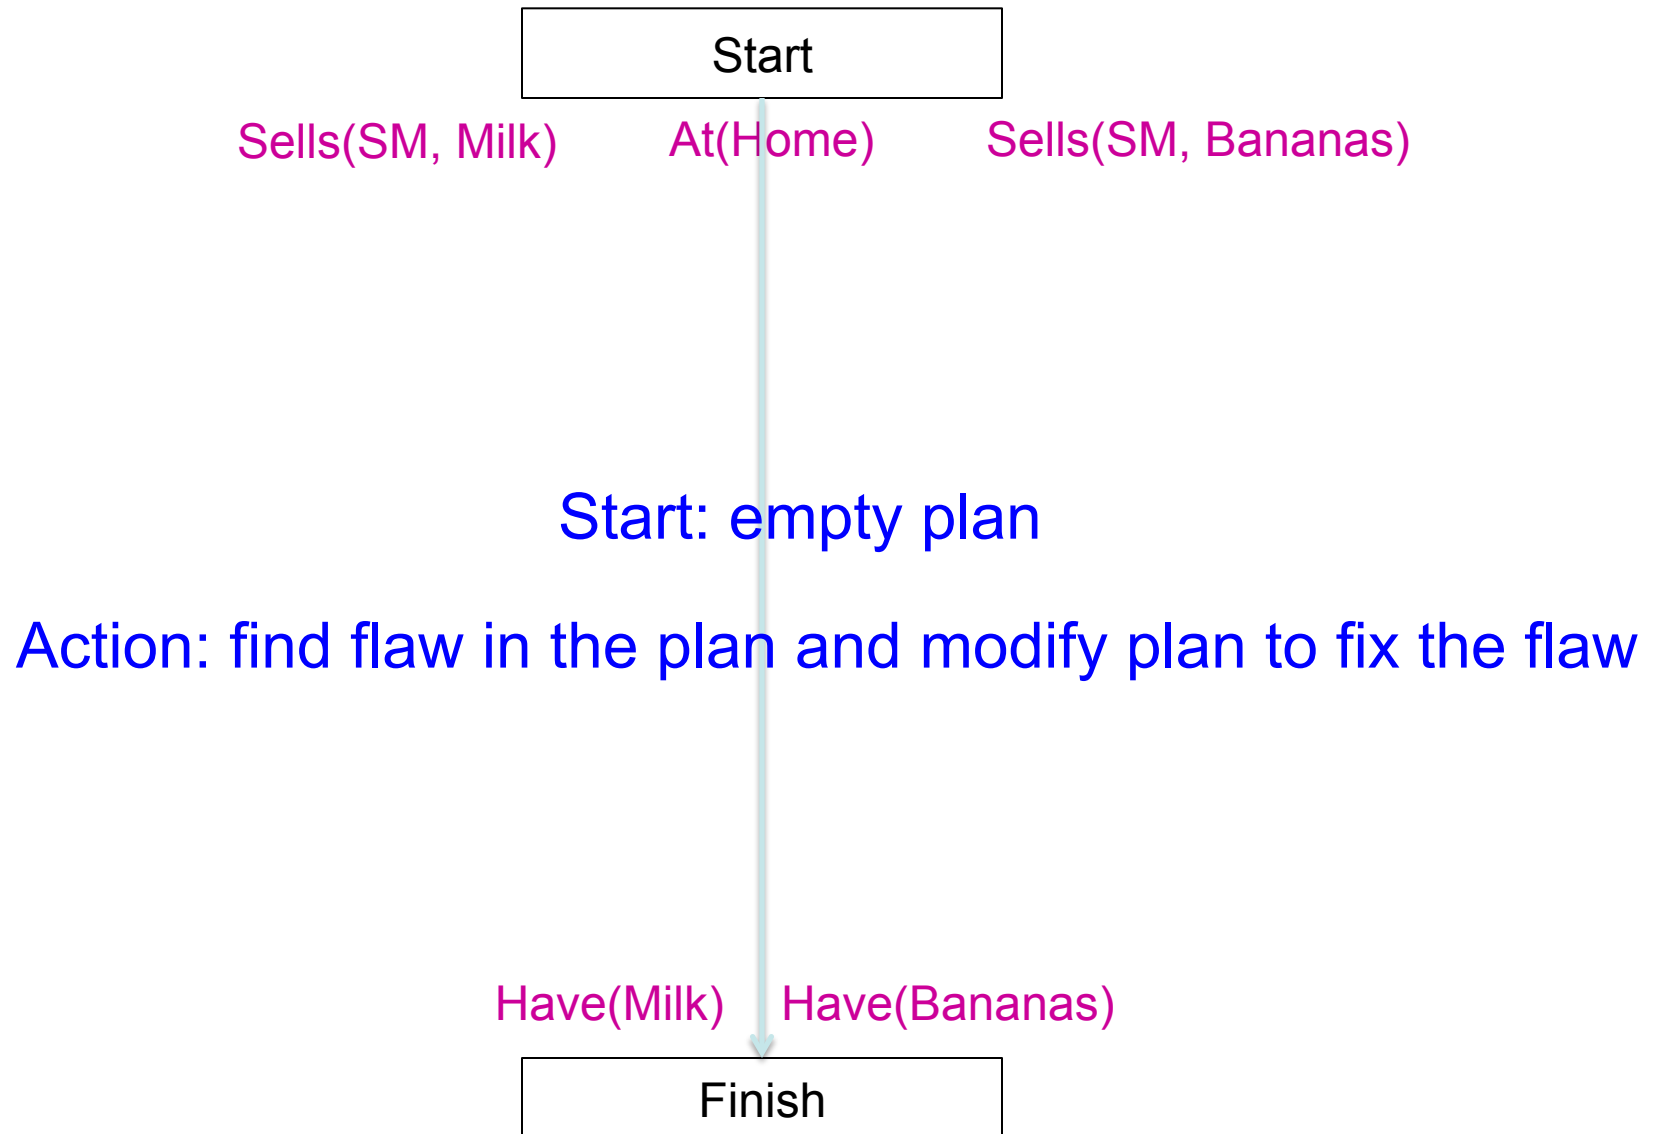
Total order (linear) plans




Partial order plan



Partial Order Planning Example



Partial Order Planning Example



Application of planning: Automated storytelling



Home

Mark Riedl



Mark Riedl, Ph.D.

Assistant Professor

School of Interactive Computing

College of Computing

Georgia Institute of Technology

GVU Center

RIM Center

Email: riedl@cc.gatech.edu

Twitter: [@mark_riedl](https://twitter.com/mark_riedl)

<https://research.cc.gatech.edu/inc/mark-riedl>

Application of planning: Automated storytelling

- Applications
 - Personalized experience in games
 - Automatically generating training scenarios (e.g., for the army)
 - Therapy for kids with autism
 - Computational study of creativity

Challenges of real-world planning

- Actions at different levels of granularity: hierarchical planning
- Resource constraints (semi-dynamic environments)
- Dynamic environments
- Stochastic or partially observable environments
- Multi-agent environments
- Example: [path planning with moving obstacles](#)