


PLANNING AND ACTING

CHAPTER 12.3 - 12.5

Outline

- ◇ The real world
- ◇ Conditional planning
- ◇ Monitoring and replanning

The real world



START

*~Flat(Spare) Intact(Spare) Off(Spare)
On(Tire1) Flat(Tire1)*

On(x) ~Flat(x)

FINISH

On(x)

Remove(x)

Off(x) ClearHub

Off(x) ClearHub

Puton(x)

On(x) ~ClearHub

Intact(x) Flat(x)

Inflate(x)

~Flat(x)

Things go wrong

Incomplete information

Unknown preconditions, e.g., *Intact(Spare)?*

Disjunctive effects, e.g., *Inflate(x)* causes

$\text{Inflated}(x) \vee \text{SlowHiss}(x) \vee \text{Burst}(x) \vee \text{BrokenPump} \vee \dots$

Incorrect information

Current state incorrect, e.g., spare NOT intact

Missing/incorrect postconditions in operators

Qualification problem:

can never finish listing all the required preconditions and possible conditional outcomes of actions

Indeterminacy in the World

Bounded indeterminacy: actions can have unpredictable effects, but the possible effects can be listed in the action description axioms

Unbounded indeterminacy: set of possible preconditions or effects either is unknown or is too large to be completely enumerated

Closely related to **qualification problem**

Solutions

Conformant or sensorless planning

Devise a plan that works regardless of state or outcome

Such plans may not exist

Conditional planning

Plan to obtain information (observation actions)

Subplan for each contingency, e.g.,

[*Check(Tire1)*, **if** *Intact(Tire1)* **then** *Inflate(Tire1)* **else** *CallAAA*

Expensive because it plans for many unlikely cases

Monitoring/Replanning

Assume normal states, outcomes


Check progress *during execution*, replan if necessary

Unanticipated outcomes may lead to failure (e.g., no AAA card)

(Really need a combination; plan for likely/serious eventualities, deal with others when they arise, as they must eventually)


Conformant planning

Search in space of **belief states** (sets of possible actual states)



Conditional planning

If the world is nondeterministic or partially observable
then percepts usually *provide information*,
i.e., *split up* the belief state



Conditional planning (con't.)

Conditional plans check (any consequence of KB +) percept

[... , **if** C **then** $Plan_A$ **else** $Plan_B$, ...]

Execution: check C against current KB, execute “then” or “else”

Conditional planning (con't.)

Need to handle nondeterminism by building into the plan conditional steps that check the state of the environment at run time, and then decide what to do.

Augment STRIPS to allow for nondeterminism:

Add **Disjunctive effects** (e.g., to model when action sometimes fails):

$Action(Left, PRECOND: AtR, EFFECT: AtL \vee AtR)$

Add **Conditional effects** (i.e., depends on state in which it's executed):

Form: **when** $\langle condition \rangle$: $\langle effect \rangle$

$Action(Suck, PRECOND:$

$EFFECT: (\mathbf{when} AtL: CleanL) \wedge (\mathbf{when} AtR: CleanR))$

Create **Conditional steps**:

if $\langle test \rangle$ **then** $plan-A$ **else** $plan-B$

Conditional planning (con't.)

Need *some* plan for *every* possible percept and action outcome

(Cf. game playing: *some* response for *every* opponent move)

(Cf. backward chaining: *some* rule such that *every* premise satisfied)

Use: AND-OR tree search (very similar to backward chaining algorithm)

Similar to game tree in minimax search

Differences: MAX and MIN nodes become OR and AND nodes

Robot takes action in “state” nodes.

Nature decides outcome at “chance” nodes.

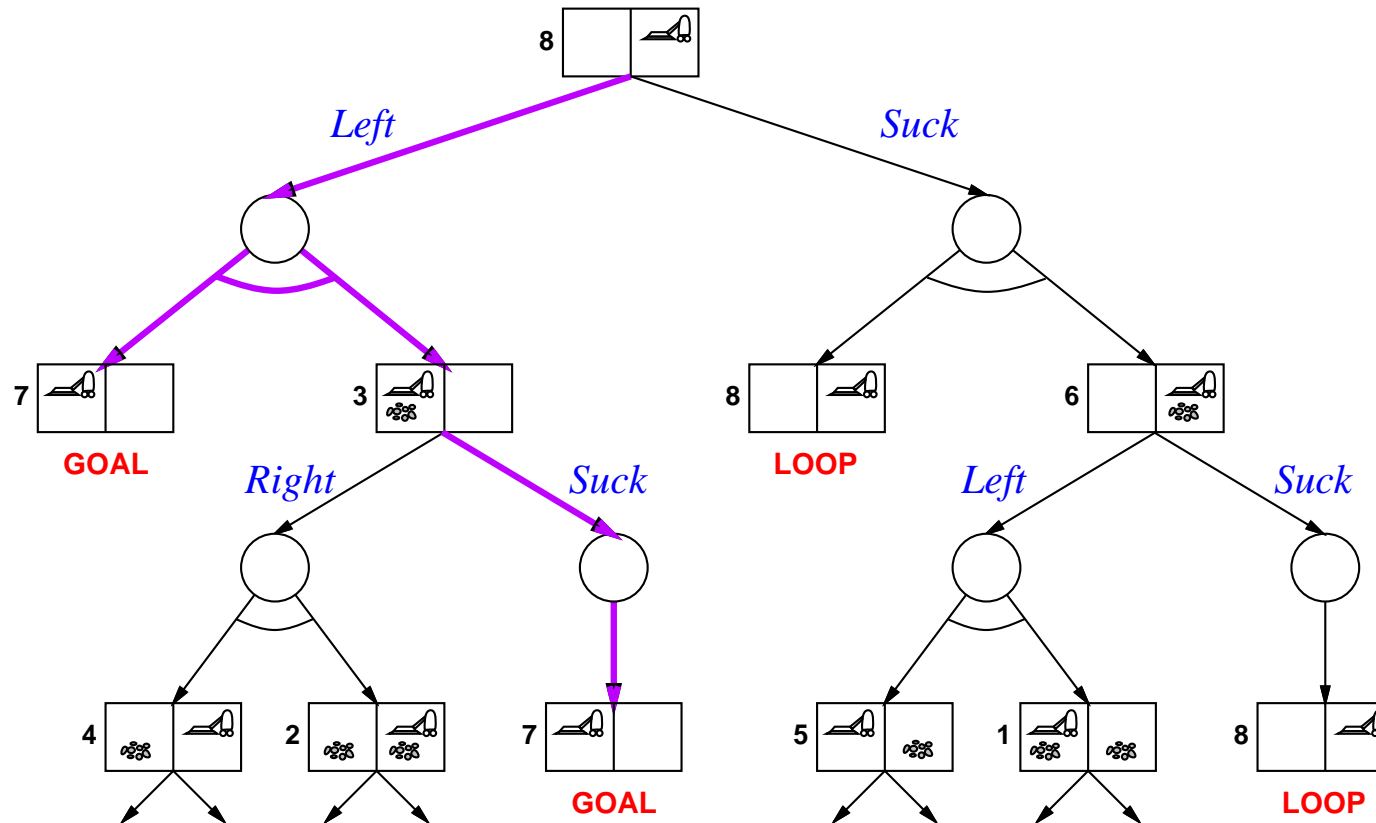
Plan needs to take *some* action at every state it reaches (i.e., OR nodes)

Plan must handle *every* outcome for the action it takes (i.e., AND nodes)

Solution is a subtree with (1) goal node at every leaf, (2) one action specified at each state node, and (3) includes every outcome branch at chance nodes.

Example: “Game Tree”, Fully Observable World

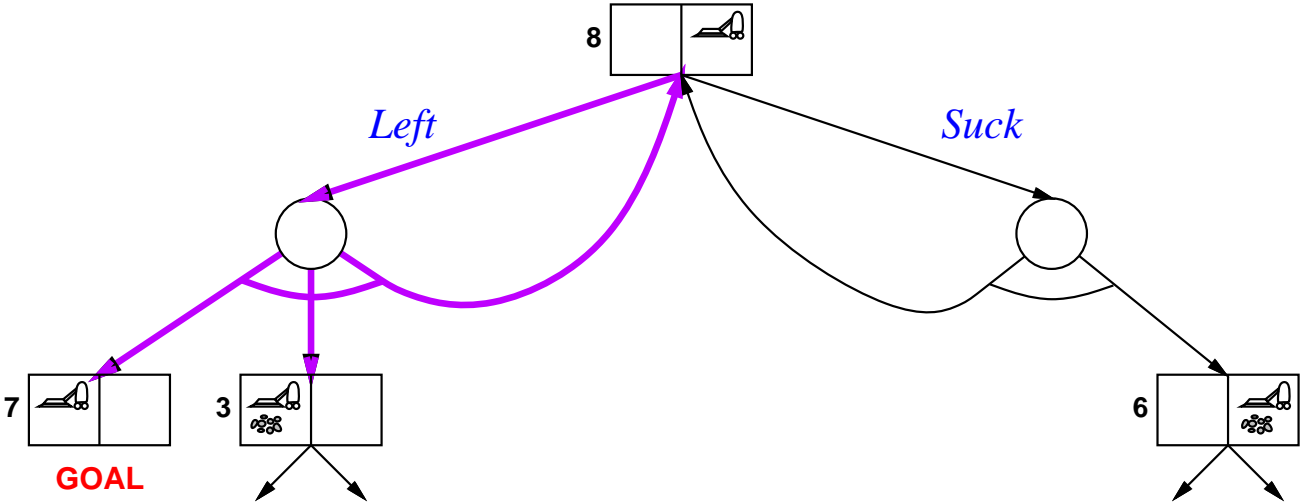
Double Murphy: sucking or arriving may dirty a clean square



Plan: $[Left, \text{if } AtL \wedge CleanL \wedge CleanR \text{ then } [] \text{ else } Suck]$

Example

Triple Murphy: also sometimes stays put instead of moving



$[L_1 : \textit{Left}, \textit{if } AtR \textit{ then } L_1 \textit{ else } [\textit{if } CleanL \textit{ then } [] \textit{ else } \textit{Suck}]]$

or $[\textit{while } AtR \textit{ do } [\textit{Left}], \textit{if } CleanL \textit{ then } [] \textit{ else } \textit{Suck}]$

“Infinite loop” but will eventually work unless action always fails

Execution Monitoring

“Failure” = preconditions of *remaining plan* not met

Preconditions of remaining plan

= all preconditions of remaining steps not achieved by remaining steps

= all causal links *crossing* current time point


On failure, resume POP to achieve open conditions from current state

IPEM (Integrated Planning, Execution, and Monitoring):

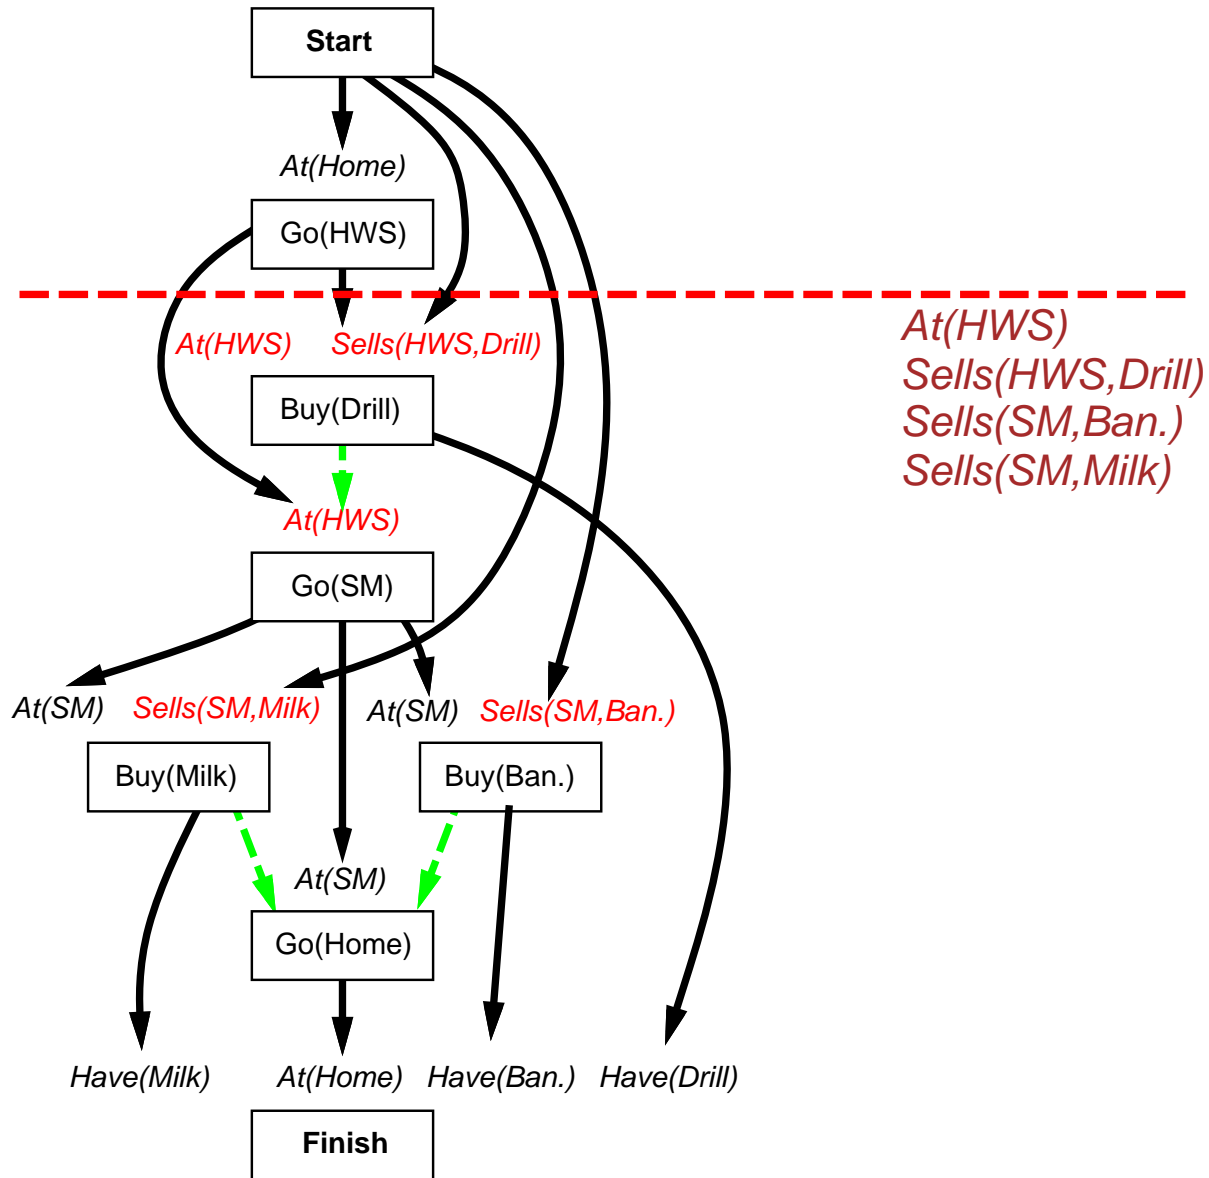
keep updating *Start* to match current state

links from actions replaced by links from *Start* when done

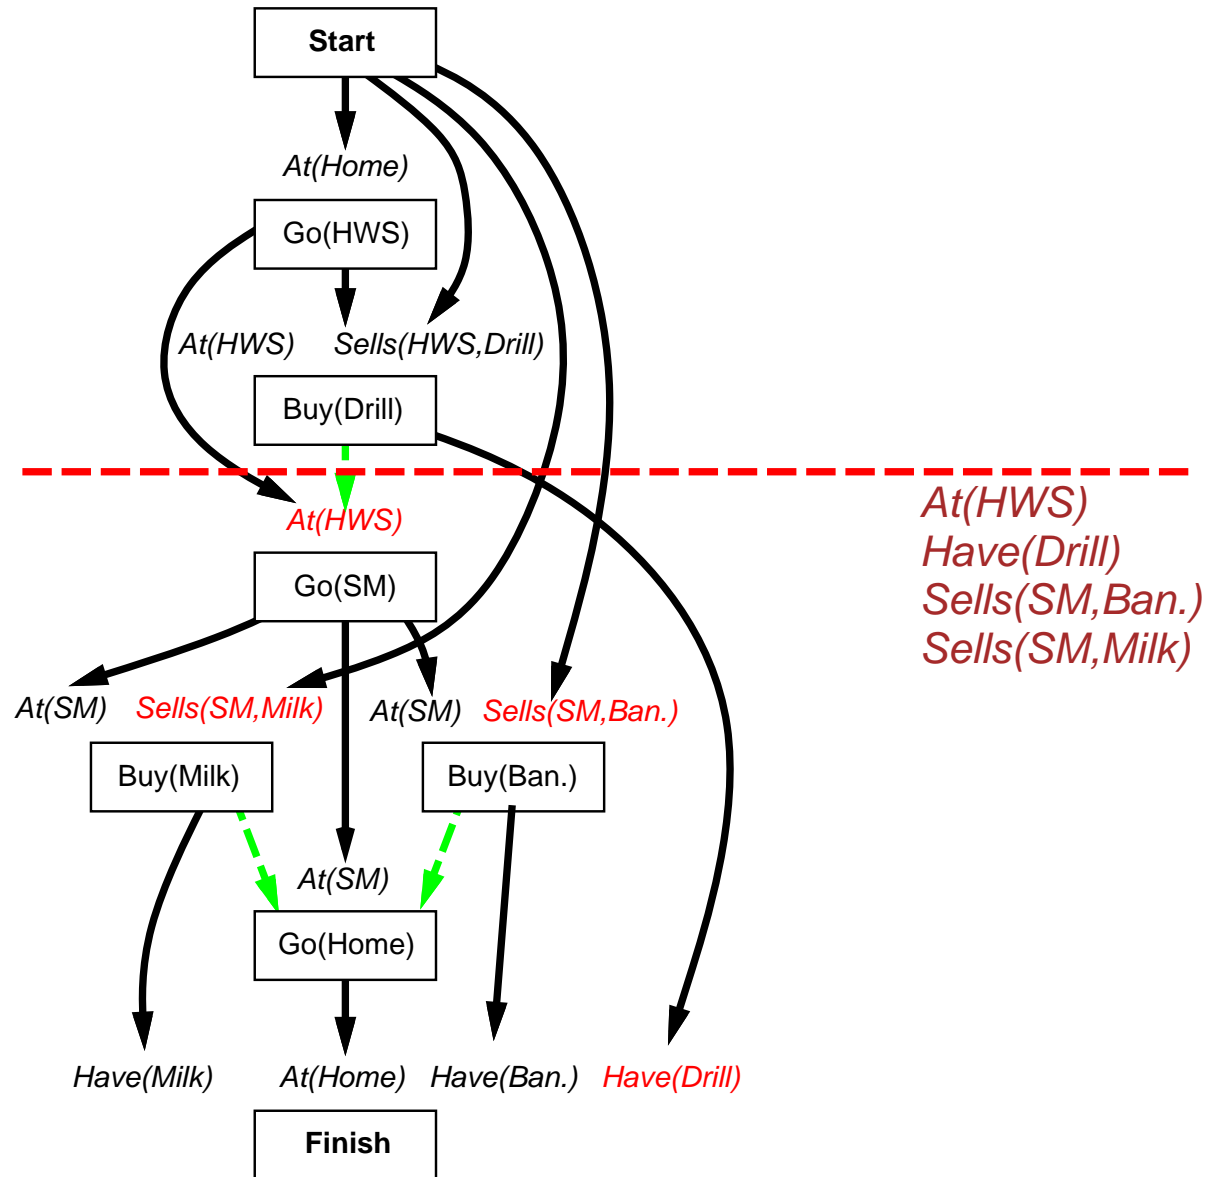
Example




Example




Example



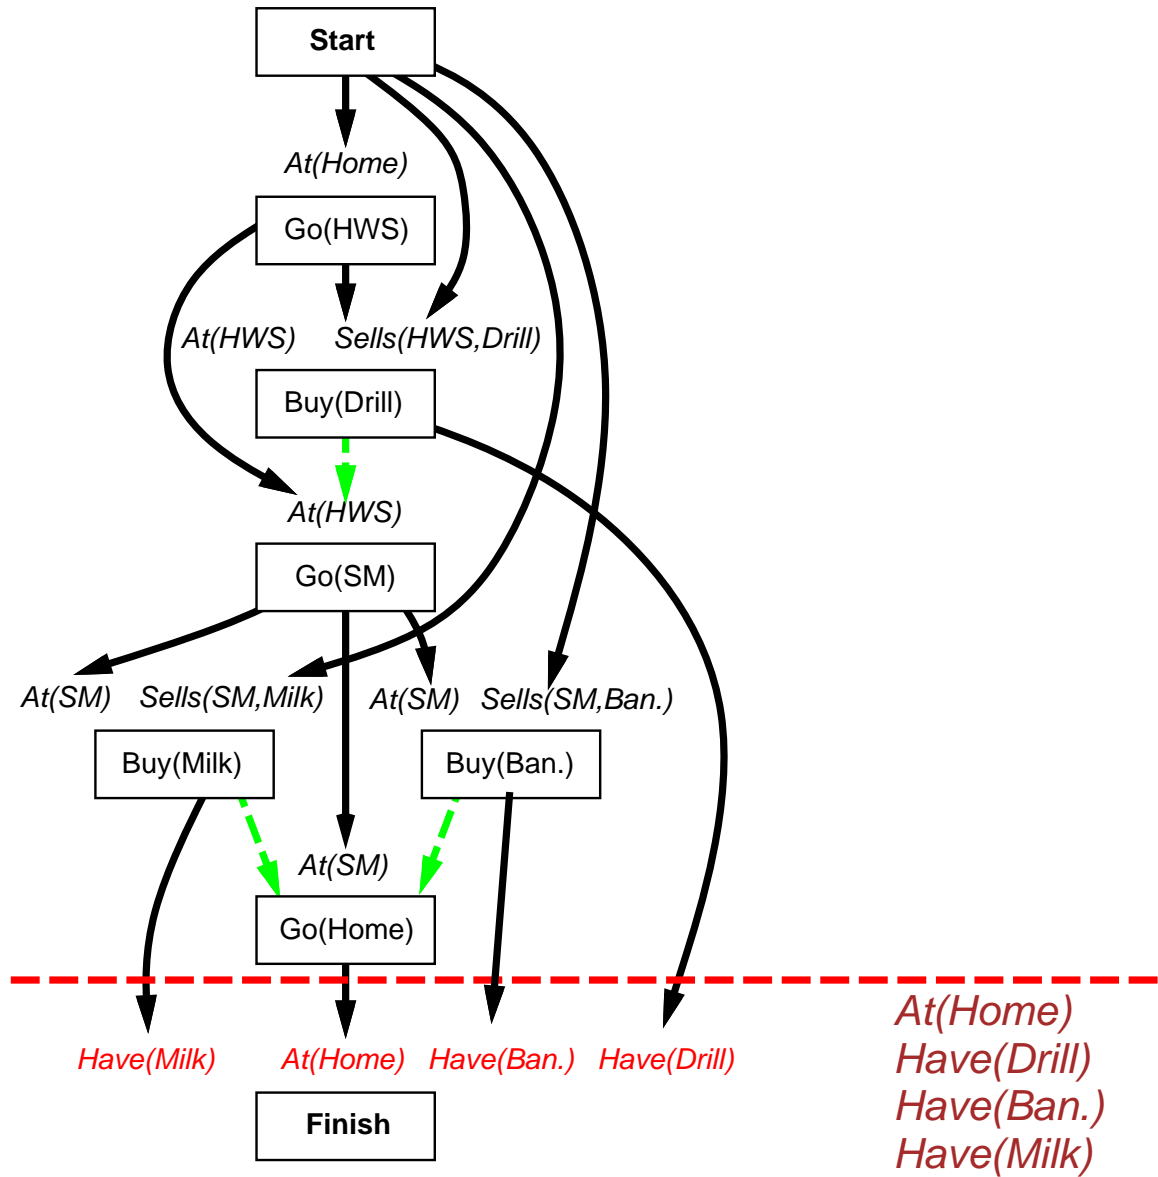
Example



Example




Example



Emergent behavior

PRECONDITIONS

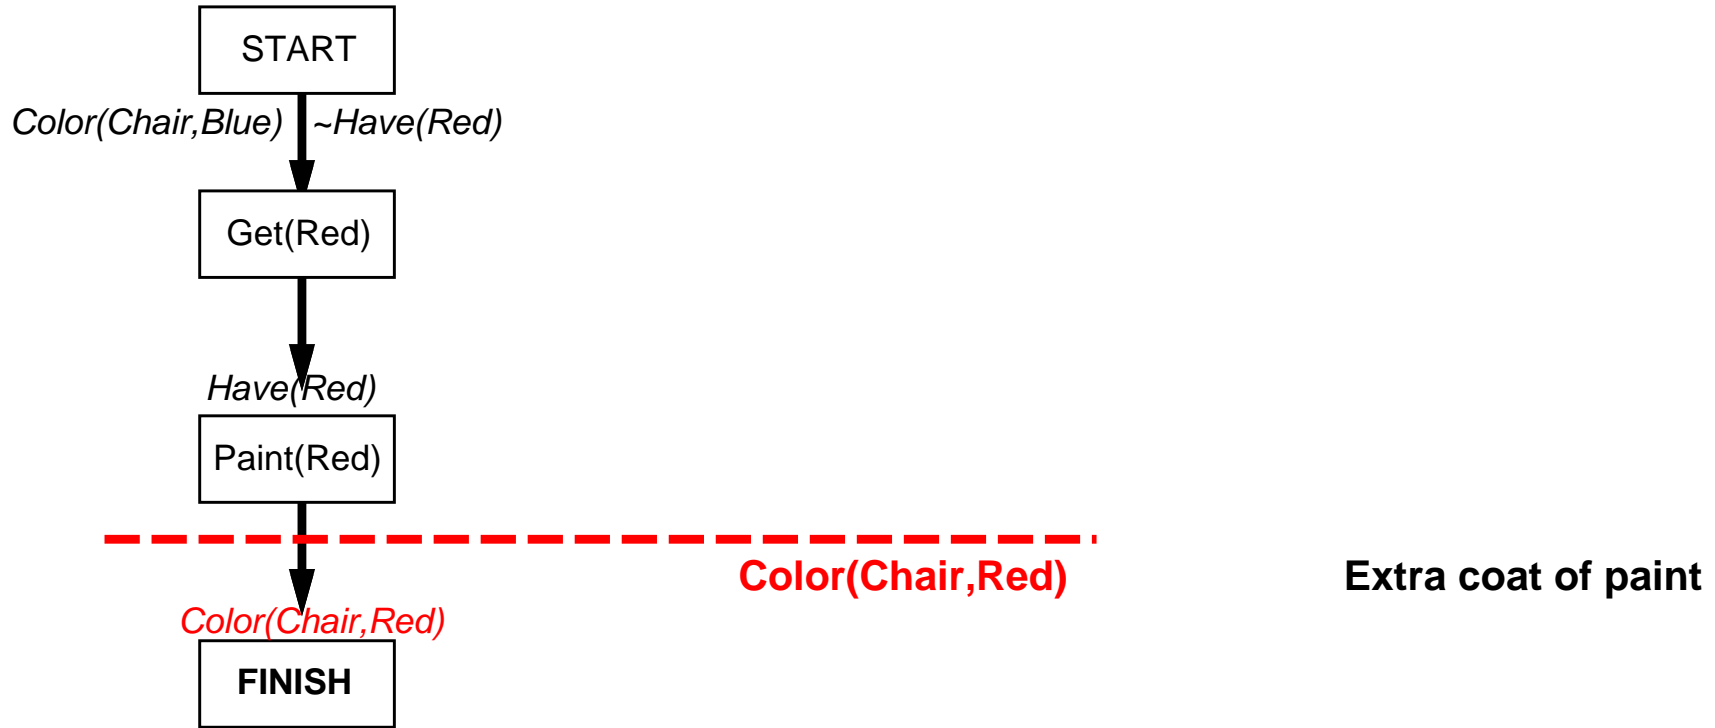
FAILURE RESPONSE



Emergent behavior

PRECONDITIONS

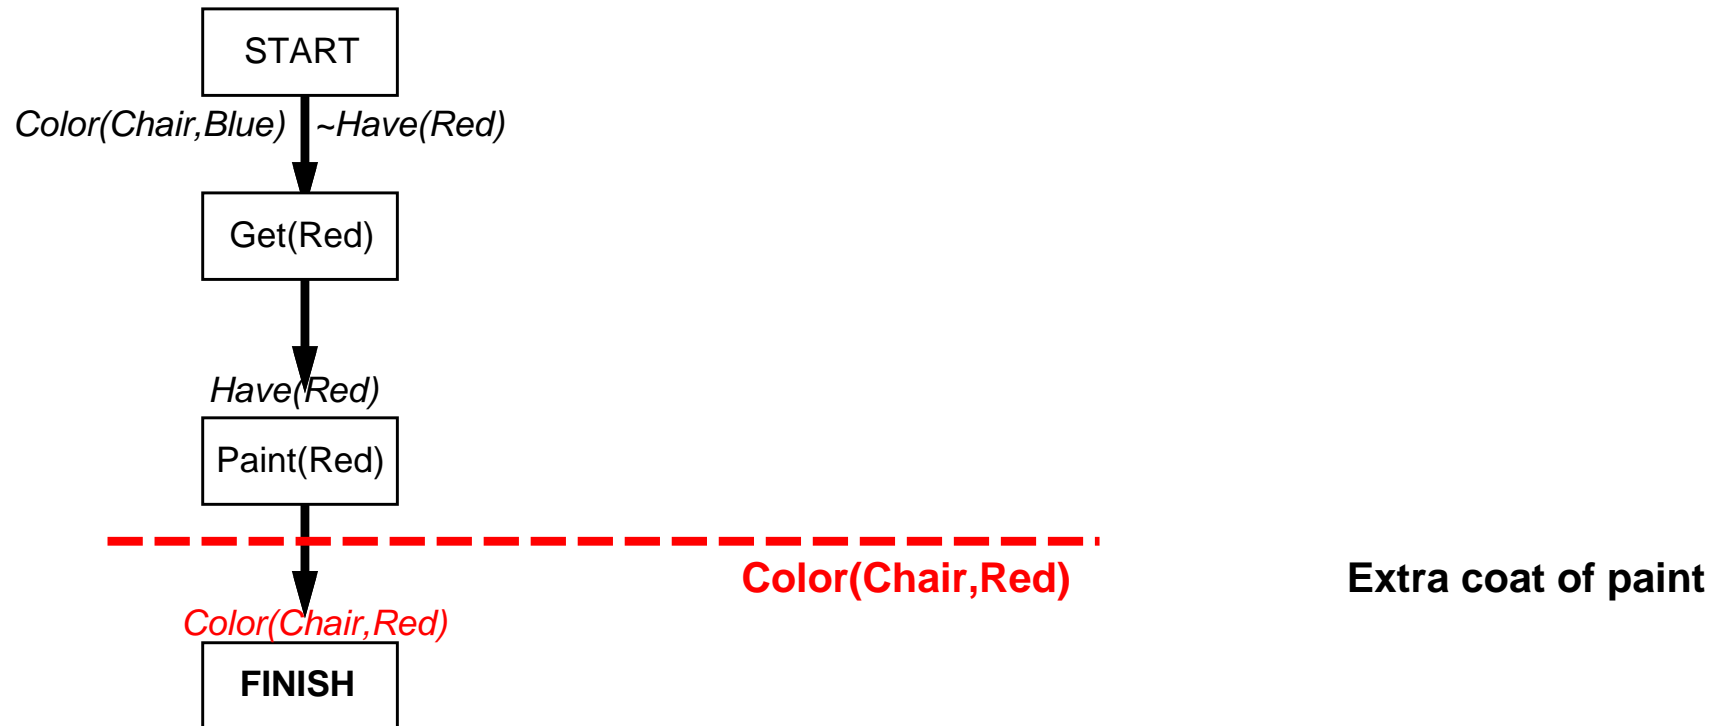
FAILURE RESPONSE



Emergent behavior

PRECONDITIONS

FAILURE RESPONSE



“Loop until success” behavior *emerges* from interaction between monitor/replan agent design and uncooperative environment

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning:

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Sensorless/Conformant planning:

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Sensorless/Conformant planning: Open can of paint and apply it to both chair and table.

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Sensorless/Conformant planning: Open can of paint and apply it to both chair and table.

Conditional planning:

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Sensorless/Conformant planning: Open can of paint and apply it to both chair and table.

Conditional planning: Sense the color of the table and chair. If same, then we're done. If not, sense labels on the paint cans; if there is a can that is the same color as one piece of furniture, then apply the paint to the other piece. Otherwise, paint both pieces with any color.

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Sensorless/Conformant planning: Open can of paint and apply it to both chair and table.

Conditional planning: Sense the color of the table and chair. If same, then we're done. If not, sense labels on the paint cans; if there is a can that is the same color as one piece of furniture, then apply the paint to the other piece. Otherwise, paint both pieces with any color.

Monitoring/replanning:

Summarizing Example

Assume: You have a chair, a table, and some cans of paint; all colors are unknown. Goal: chair and table have same color.

How would each of the following handle this problem?

Classical planning: Can't handle it, because initial state isn't fully specified.

Sensorless/Conformant planning: Open can of paint and apply it to both chair and table.

Conditional planning: Sense the color of the table and chair. If same, then we're done. If not, sense labels on the paint cans; if there is a can that is the same color as one piece of furniture, then apply the paint to the other piece. Otherwise, paint both pieces with any color.

Monitoring/replanning: Similar to conditional planner, but perhaps with fewer branches at first, which are filled in as needed at runtime. Also, would check for unexpected outcomes (e.g., missed a spot in painting, so repaint)

Summary

- ◇ Incomplete info: use conditional plans, conformant planning (can use belief states)
- ◇ Incorrect info: use execution monitoring and replanning