



Automated Planning

Domain-Configurable Planning: Planning with Control Formulas

Jonas Kvarnström

Department of Computer and Information Science

Linköping University

Domain-Configurable: Prioritization and Pruning

Two Kinds of Search Guidance (1)

Prioritization: Which part of a search tree should be visited first?
Could use heuristic functions, could use other methods...

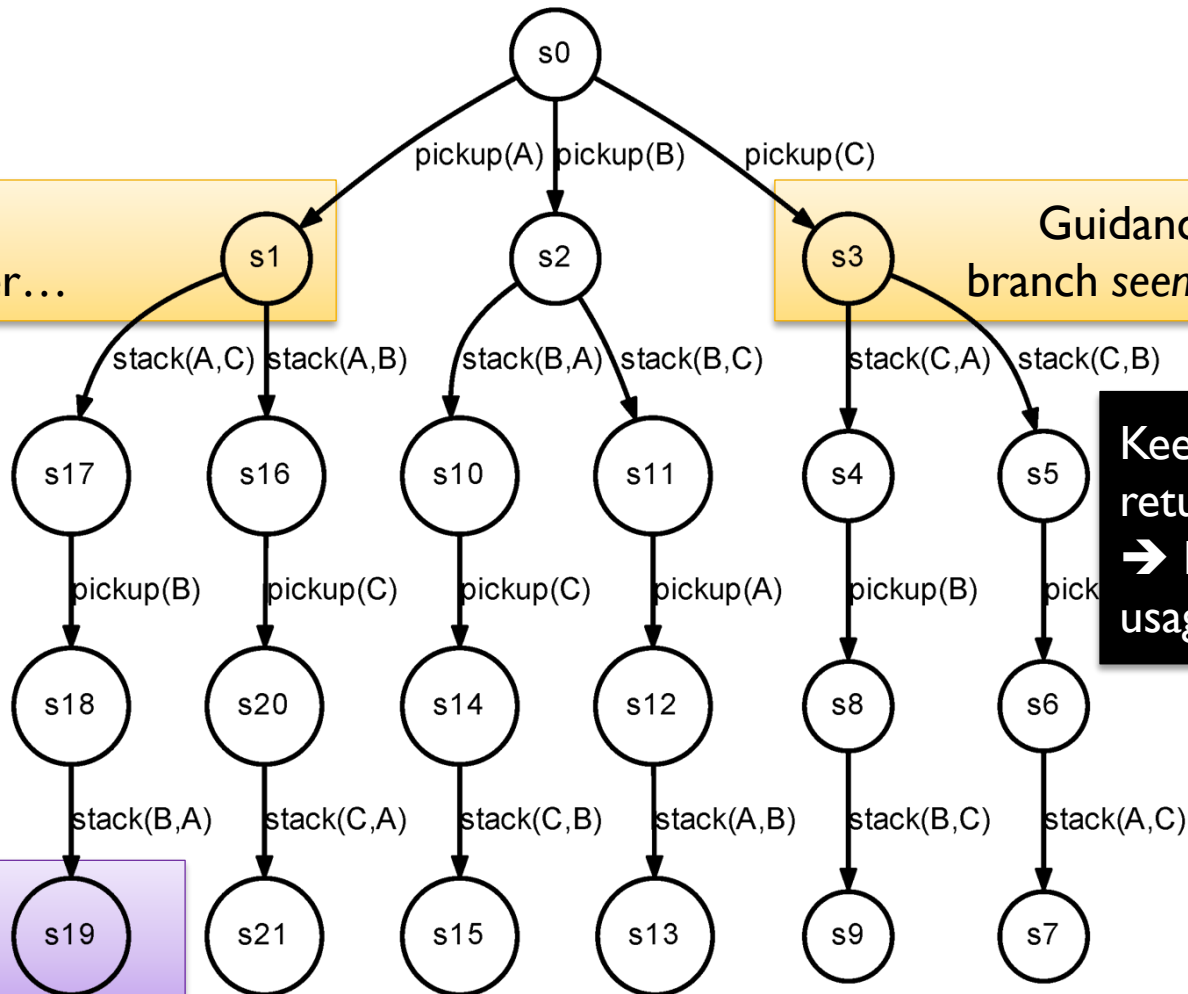
Guidance says: This branch seems better...

Go this way first

Guidance says: This branch seems worse...

Keep this, maybe return later
→ high memory usage

Goal:
B on C on A

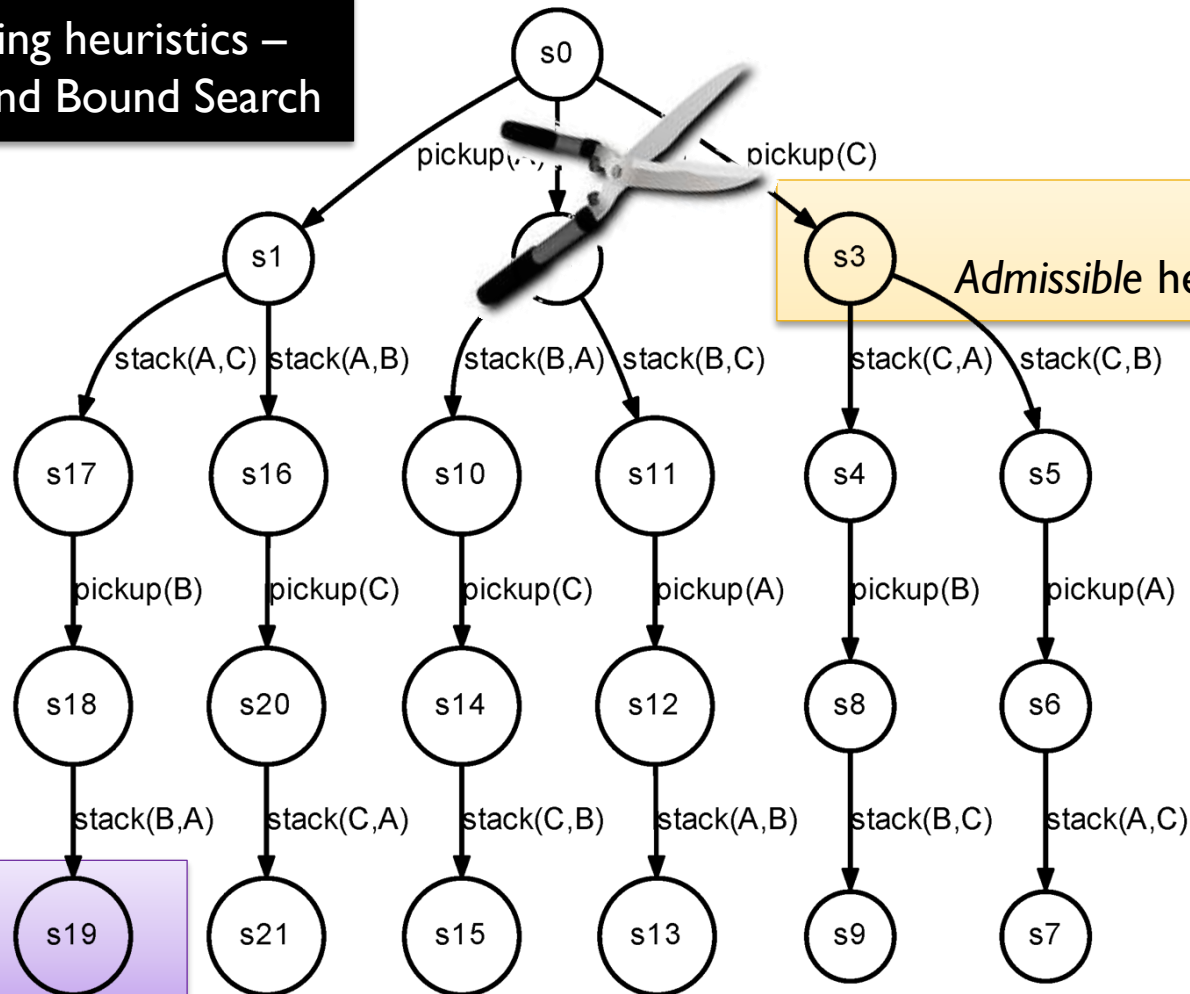


- Properties of prioritization:
 - We can always return to a node later
 - No need to be absolutely certain of your priorities
- This is why many domain-independent heuristics work well
 - Provide *reasonable* advice in *most* cases

Two Kinds of Search Guidance (2)

Pruning: Which part of a search tree are definitely useless?
Prune them!

Can be done using heuristics –
example: Branch and Bound Search



Cost $g(n)=1$
Admissible heuristic $h(n)=4$
Any solution
below s_3
will cost
at least 5

Prune!
Never
consider the
node or its
descendants
again...


We found a
solution of cost 4

- Can we prune when we search for the **first** solution?
 - A single mistake may *remove all paths to solutions*
 - → Difficult to find good *domain-independent* pruning criteria

Example: Emergency Services Logistics

- Emergency Services Logistics

- Goal: `at(crate1, loc1), at(crate2, loc2), at(crate3, loc3)`
- Now: `at(crate1, loc1), at(crate2, loc2)`




- Picking up crate1 again is **physically possible**
- It “destroys” `at(crate1, loc1)`, which is a goal – **obviously stupid!**

The **branch** beginning with `pickup(crate1)` could be **pruned** from the tree!
How do we **detect** this in a domain-independent way?

Example: Towers of Hanoi

Should we always prevent the destruction of achieved goals?


- Goal: on(1,2), on(2,3), on(3,4), on(4,5), on(5,6), on(6,7)
- Now: on(1,2), on(2,3), on(3,4), on(5,6), on(6,7)




- Moving disk 1 to the third peg is possible but “destroys” a goal fact: on(1,2)
 - Is this also obviously stupid?
 - No, it is necessary! Disk 1 is blocking us from moving disk 4...

→ Heuristics may or may not detect:

Picking up crate 1 is bad




Moving disk 1 is good



Will generally depend on the entire state
+ which alternative states exist,
not just the fact that you "destroy goal achievement"

Might delay investigating either alternative for a while,
return to try this later

- With a **domain-configurable planner**:
 - We *could* provide *domain-specific heuristics*
 - Strongly discouraging the destruction of goals in Emergency Services Logistics
 - Would keep the *option* to investigate such actions later (not necessary!)
 - We can directly provide stronger domain-specific **pruning criteria**



Planning with Control Formulas


- **Control formulas**: One way of specifying when to prune
 - Motivation
 - Examples
 - Formalism
 - Evaluation of control formulas

Simplest control information: Precondition Control

- **operator** `pickup(robot, crate, location)`

- **precond:**

- `at(robot, location), at(crate, location)`
- `handempty(robot)`
- ...and **the goal doesn't require that crate should end up at location!**



How to express this, given that the goal requires `at(crate, loc)`?

- Alternative 1: **New predicate** "`destination(crate, loc)`"

- *Duplicates* the information already specified in the goal
- **precond:** `¬destination(crate, location)`

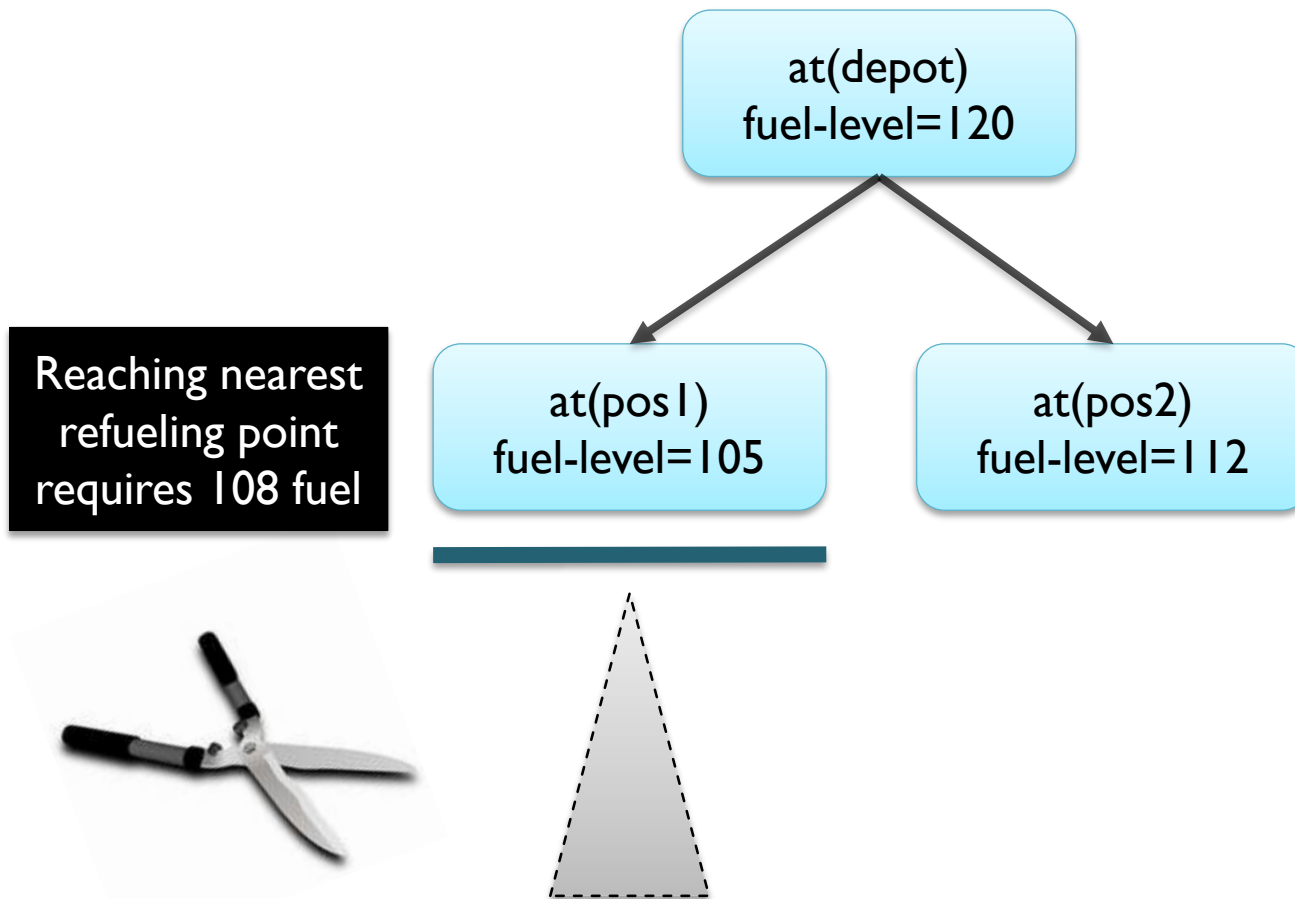

Supported by
all planners

- Alternative 2: **New language extension** "`goal(ϕ)`"


- Evaluated in the set of goal states, not in the current state
- **precond:** `¬goal(at(crate, location))`

Requires
extensions, but
more convenient

- A UAV should never be where it **can't reach** a refueling point
 - Can't possibly extend such plans into solutions



- A UAV should never be where it **can't reach** a refueling point
 - If this happens in a plan, we can't possibly extend it into a solution satisfying the goal
- How to express this?



Using preconditions again?

Must be verified for **every** action:
fly, scan-area, take-off, ...

Must be checked even when
the UAV is idle, hovering

Inconvenient!

Using state constraints?

Defined **once**,
applied to **every generated state**

```

$$\forall(u: \text{uav}) [$$

$$\quad \exists(rp: \text{refueling-point}) [$$

$$\quad \quad \text{dist}(u, rp) * \text{fuel-usage}(u) < \text{fuel-avail}(u)$$

$$\quad ]$$

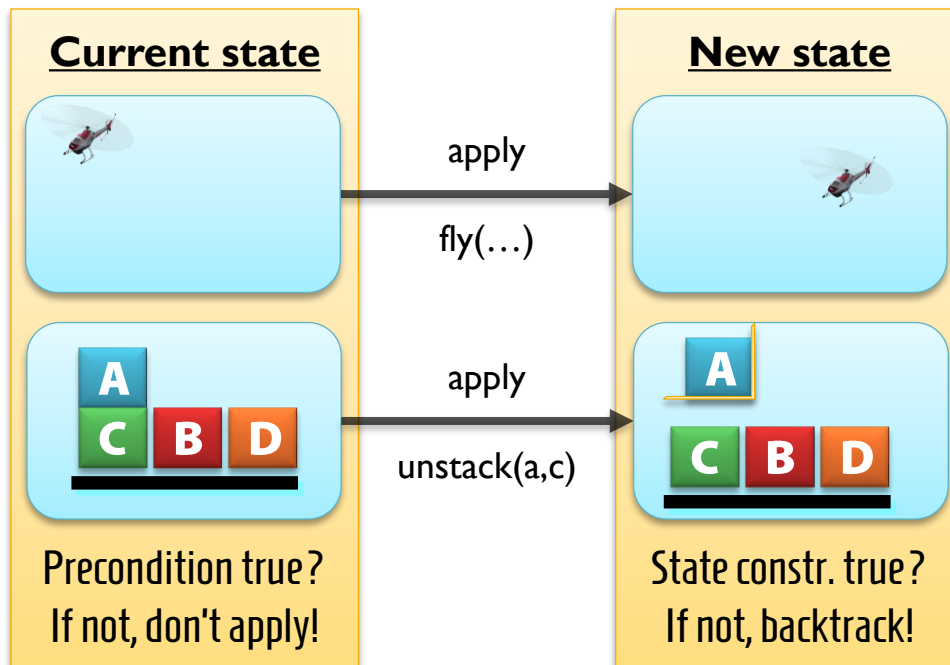
$$]$$

```

Comparatively simple extension!

Testing State Constraints

- Testing such state constraints is simple
 - Apply an action → new state is generated
 - Formula false in that state → Prune!
 - Similar to preconditions
 - But tested in the state after an action is applied, not before!



Temporal Conditions (1)




- A package on a carrier should **remain there** until it reaches its destination
 - For any plan π where we move it prematurely, there is a more efficient plan π' where we don't

How to express this as a single formula?



Temporal Conditions (2)

- “A package on a carrier should **remain** there **until** it reaches its destination”

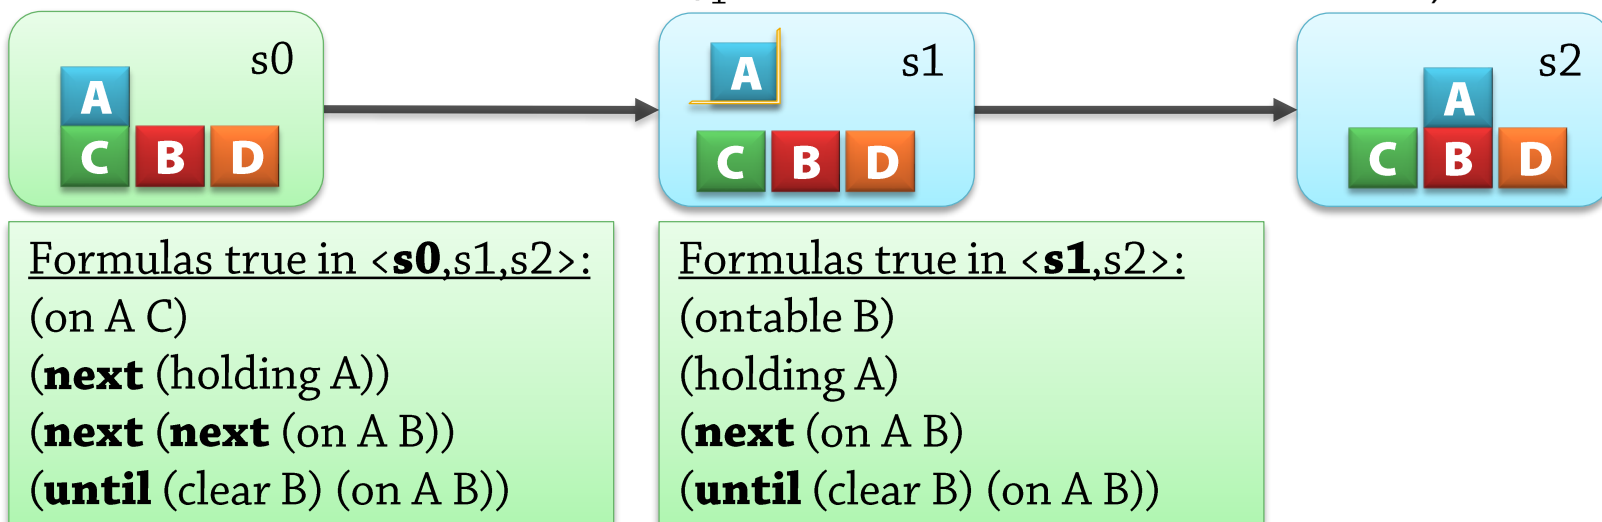


We need a formula constraining an entire **state sequence**, not a single state!

In planning, this is called a **control formula** or **control rule**

We need to extend the logical language!

- One possibility: Use **Linear Temporal Logic** (as in TLplan)
 - All formulas evaluated relative to a *state sequence* and a *current state*
 - Assuming that f is a formula:
 - $\bigcirc f$ – (**next** f) f is true in the next state
 - $\diamond f$ – (**eventually** f) f is true either now **or** in some future state
 - $\square f$ – (**always** f) f is true now **and** in all future states
 - $f_1 \cup f_2$ – (**until** $f_1 f_2$) (f_2 is true now), **or** (f_2 is true in some future state s' and f_1 is true in all states until/before then)



- “A package on a carrier should remain there until it reaches its destination”

(forall (?var) (type-predicate ?var) ϕ):
 $\forall v. \text{type-predicate}(v) \rightarrow \phi$
For all values of ?var that satisfy type-predicate, ϕ must be true

- (always

(**forall** (?c) (carrier ?c)

;; For all carriers

(**forall** (?p) (package ?p)

;; For all packages

(**implies**

(on-carrier ?p ?c)

;; If the package is on the carrier

(**until** (on-carrier ?p ?c)

;; ...then it remains on the carrier

(**exists** (?loc)

;; until there exists a location

(at ?p ?loc)

;; where it is, and

(**goal** (at ?p ?loc))))


;; where the goal says it should be

))))

- Should be true starting in the initial state

Control Formula

- (always**
 (forall (?c) (carrier ?c) ;; For all carriers
 (forall (?p) (package ?p) ;; For all packages
 (implies
 (on-carrier ?p ?c) ;; If the package is on the carrier
 (until (on-carrier ?p ?c) ;; ...then it remains on the carrier
 (exists (?loc) ;; until there exists a location
 (at ?p ?loc) ;; where it is, and
 (goal (at ?p ?loc)))) ;; where the goal says it should be



always
→
time 0

p1 on c1 here

Remains on c1 until at its destination

always
→
time 1

p1 on c1 here

p2 on c2 here

Until at its dest

Until at its dest

always
time 2

p2 on c2 here ...

Finding Control Formulas

- How do we come up with good control rules?
 - Good starting point: "Don't be stupid!"
 - Trace the search process – suppose the planner tries this:



- Placing F on top of B is stupid, because we'll have to remove it later
 - Would have been better to put F on the table!
- Conclusion: Should not extend a good tower the wrong way
 - *Good tower*: a tower of blocks that will never need to be moved

Blocks World Example (continued)

- Rule 1: Every goodtower must always remain a goodtower
 - **(forall** (?x) (clear ?x) ;; For all blocks that are clear (at the top of a tower)
 (implies
 (goodtower ?x) ;; If the tower is good (no need to move any blocks)
 (next (or ;; ...then in the next state, either:
 (clear ?x) ;; ?x remains clear (didn't extend the tower)
 (exists (?y) (on ?y ?x) ;; or there is a block ?y which is on ?x
 (goodtower ?y)) ;; which is a goodtower
)))

s0

s1

s2

s3

goodtower(x)? → clear(x) or
goodtower(y)

What about the rest?

Blocks World Example (continued)

■ Rule 1, second attempt:

■ **(always**

(forall (?x) (clear ?x) ;; For all blocks that are clear (at the top of a tower)

(implies

(goodtower ?x) ;; If the tower is good (no need to move any blocks)

(next (or ;; ...then in the next state, either:

(clear ?x) ;; ?x remains clear (didn't extend the tower)

(exists (?y) (on ?y ?x) ;; or there is a block ?y which is on ?x

(goodtower ?y)) ;; which is a goodtower

))))

s0

s1

s2

s3

goodtower(x)? → clear(x) or
goodtower(y)

goodtower(x)? → clear(x) or
goodtower(y)

goodtower(x)? → clear(x) or

Supporting Predicates

- Some planners allow us to **define** a predicate recursively

- goodtowerbelow**(x) means we **will not have to move** x

$$\text{goodtowerbelow}(x) \Leftrightarrow [\text{ontable}(x) \wedge \neg \exists [y: \mathbf{GOAL}(\text{on}(x,y))]]$$

∨

$$\begin{aligned} & \exists [y: \text{on}(x,y)] \{ \\ & \quad \neg \mathbf{GOAL}(\text{ontable}(x)) \wedge \\ & \quad \neg \mathbf{GOAL}(\text{holding}(y)) \wedge \\ & \quad \neg \mathbf{GOAL}(\text{clear}(y)) \wedge \\ & \quad \forall [z: \mathbf{GOAL}(\text{on}(x,z))] (z = y) \wedge \\ & \quad \forall [z: \mathbf{GOAL}(\text{on}(z,y))] (z = x) \wedge \\ & \quad \text{goodtowerbelow}(y) \\ & \} \end{aligned}$$

X is on the table,
and shouldn't be on anything else

X is on something else

Shouldn't be on the table,
shouldn't be holding it,
shouldn't be clear

If x should be on z, then it is (z is y)

If z should be on y, then it is (z is x)

The remainder of the tower is also good



goodtowerbelow: B, C, H



Supporting Predicates

- **goodtower**(x) means x is the block at the top of a good tower
 - $goodtower(x) \Leftrightarrow clear(x) \wedge \neg GOAL(holding(x)) \wedge goodtowerbelow(x)$
- **badtower**(x) means x is the top of a tower that isn't good
 - $badtower(x) \Leftrightarrow clear(x) \wedge \neg goodtower(x)$



goodtower: B
goodtowerbelow: B, C, H
badtower: G, E
(neither: D, A)



■ Step 2: Is this stupid?



- Placing **F on top of E** is stupid, because we have to move E later...
 - Would have been better to put F on the table!
 - But E was not a goodtower, so the previous rule didn't detect the problem
- Never put anything on a badtower!
 - (always
 - (forall (?x) (clear ?x) ;; For all blocks at the top of a tower
 - (implies
 - (badtower ?x) ;; If the tower is bad (must be dismantled)
 - (next (not (exists (?y) (on ?y ?x)))))) ;; Don't extend it!

■ Step 3: Is this stupid?



■ Picking up F is stupid!

- It is on the table, so we can wait until its destination is ready:



■ (always

(**forall** (?x) (clear ?x) ;; For all blocks at the top of a tower

(**implies**

(**and** (ontable ?x)

(exists (?y) (goal (on ?x ?y)) (not (goodtower ?y))))

(**next** (not (holding ?x))))))

Pruning using Control Formulas

Pruning using Control Formulas



- How do we decide when to prune the search tree?
 - Obvious idea:
 - Take the state sequence corresponding to the current action sequence
 - Evaluate the formula over that sequence
 - If it is false: Prune / backtrack!

■ Problem:

■ (always

(**forall** (?c) (carrier ?c)

;; For all carriers

(**forall** (?p) (package ?c)

;; For all packages

(**implies**

(on-carrier ?p ?c)

;; If the package is on the carrier

(**until** (on-carrier ?p ?c)

;; ...then it remains on the carrier

(**exists** (?loc)

;; until there exists a location

(at ?p ?loc)

;; where it is, and

(**goal** (at ?p ?loc))))

;; where the goal says it should be

))))

No package on a carrier
in the initial state:
Everything is OK



s0

**"Every boat I own
is a billion-dollar yacht
(because I own zero boats)"**

■ Problem:

■ (always

(**forall** (?c) (carrier ?c)

;; For all carriers

(**forall** (?p) (package ?c)

;; For all packages

(**implies**

(on-carrier ?p ?c)

;; If the package is on the carrier

(**until** (on-carrier ?p ?c)

;; ...then it remains on the carrier

(**exists** (?loc)

;; until there exists a location

(at ?p ?loc)

;; where it is, and

(**goal** (at ?p ?loc))))

;; where the goal says it should be

))))

When we add an action
placing a package
on a carrier...

...there is no future state
where the package is
at its destination!

s0


s1
(on-carrier p4 c4)

The formula is violated,
but only because the solution is not *complete* yet!
We must be allowed to continue,
generating new states...

Evaluation 3: What's Wrong?

- We had an **obvious** idea:
 - Take the state sequence corresponding to the current plan
 - Evaluate the formula over that sequence
 - If it is false: Prune / backtrack!
- This is actually **wrong**!
 - Formulas should hold in the state sequence of the **solution**
 - But they don't have to hold in every **intermediate** action sequence...

■ Analysis:



We have applied some actions, yielding a sequence of states


We intend to generate additional actions and states, but right now we don't know which ones

The control formula should be satisfied by the entire state sequence corresponding to a solution

We only know some of those states

Should only backtrack if we can prove that you can't find additional states so that the control formula becomes true

■ Analysis 2:



The control formula should be satisfied
by the entire state sequence corresponding to a solution

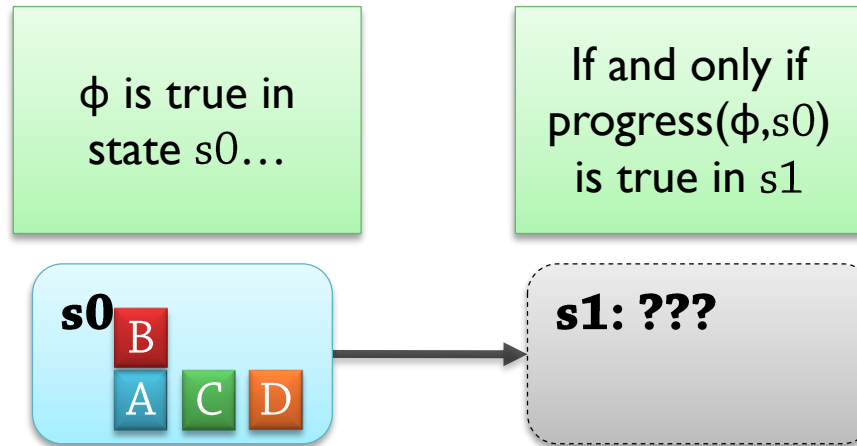
Evaluate those **parts** of
the formula that refer to
known states

Leave other parts of the formula
to be evaluated later

If the result can be proven to be **FALSE**, then backtrack

Progressing Temporal Formulas (1)

- We use formula progression
 - We progress a formula Φ through a single state s at a time
 - First the initial state, then each state generated by adding an action
 - The result is a new formula
 - Containing conditions that we must "postpone", evaluate starting in the next state



Progressing Temporal Formulas (2)

- More intuitions?

- Suppose you are reading a book. A *page* is analogous to a *state*.

"The figure appears 5 pages from here"
is true on page 7



"The figure appears 4 pages from here"
is true on page 8

"Fig 1 appears 3 pages from here
and fig 2 appears 5 pages from here"
is true on page 7



"Fig 1 appears 2 pages from here
and fig 2 appears 4 pages from here"
is true on page 8

"Starting where I am now,
there's a figure on every page"
is true on page 7



There actually is a figure on page 7, and
"Starting where I am now,
there's a figure on every page"
is true on page 8

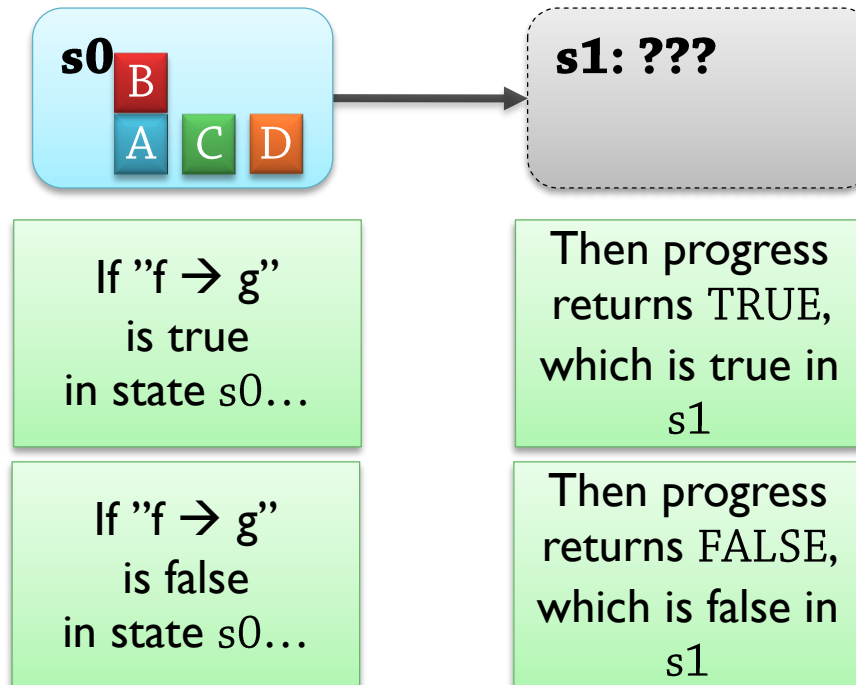
"If there's a figure on this page,
there is no figure on the next page
[otherwise I don't care]"
is true on page 7



Check if there is a figure on this page.
If so: $f = \text{"There's no figure on this page"}$.
Otherwise: $f = \text{"true"}$.
Check whether f is true on page 8.

Progressing Temporal Formulas (3)

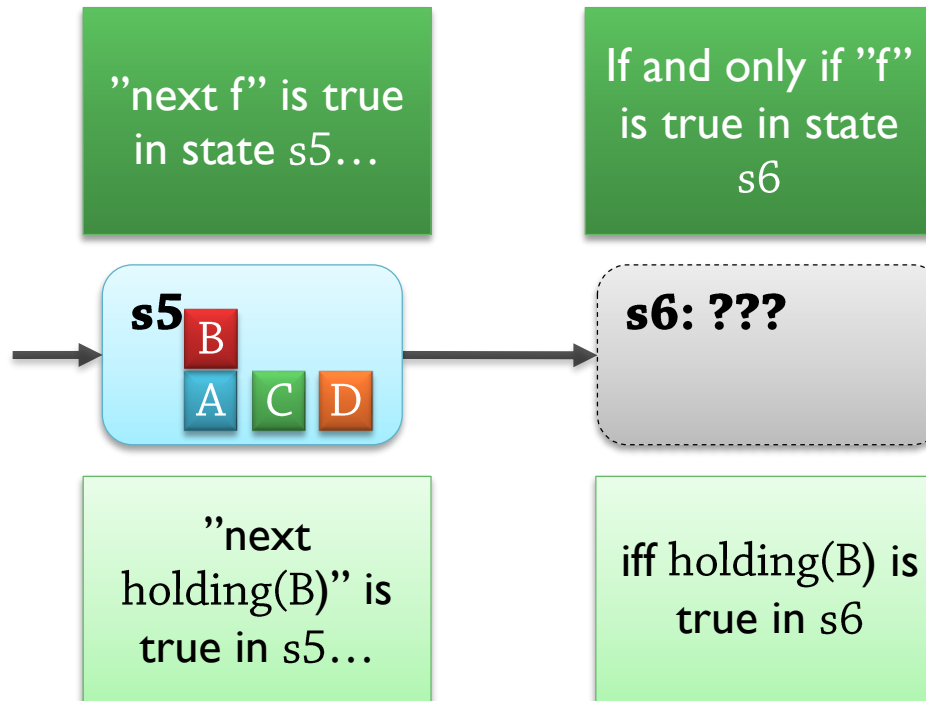
- Base case: Formulas **without** temporal operators (“ $\text{on}(A,B) \rightarrow \text{on}(C,D)$ ”)
 - Must be true **here**, in **this** state
 - $\text{progress}(\Phi, s) = \text{TRUE}$ if Φ holds in s (we already know how to test this)
 - $\text{progress}(\Phi, s) = \text{FALSE}$ otherwise



Progressing Temporal Formulas (4)

■ Simple case: next

- $\text{progress}(\text{next } f, s) = f$
 - Because "next f " is true in this state iff f is true in the next state
 - This is by definition what $\text{progress}()$ should return!



Additional cases are discussed in the book (always, eventually, until, ...)

Progression in Depth First Search

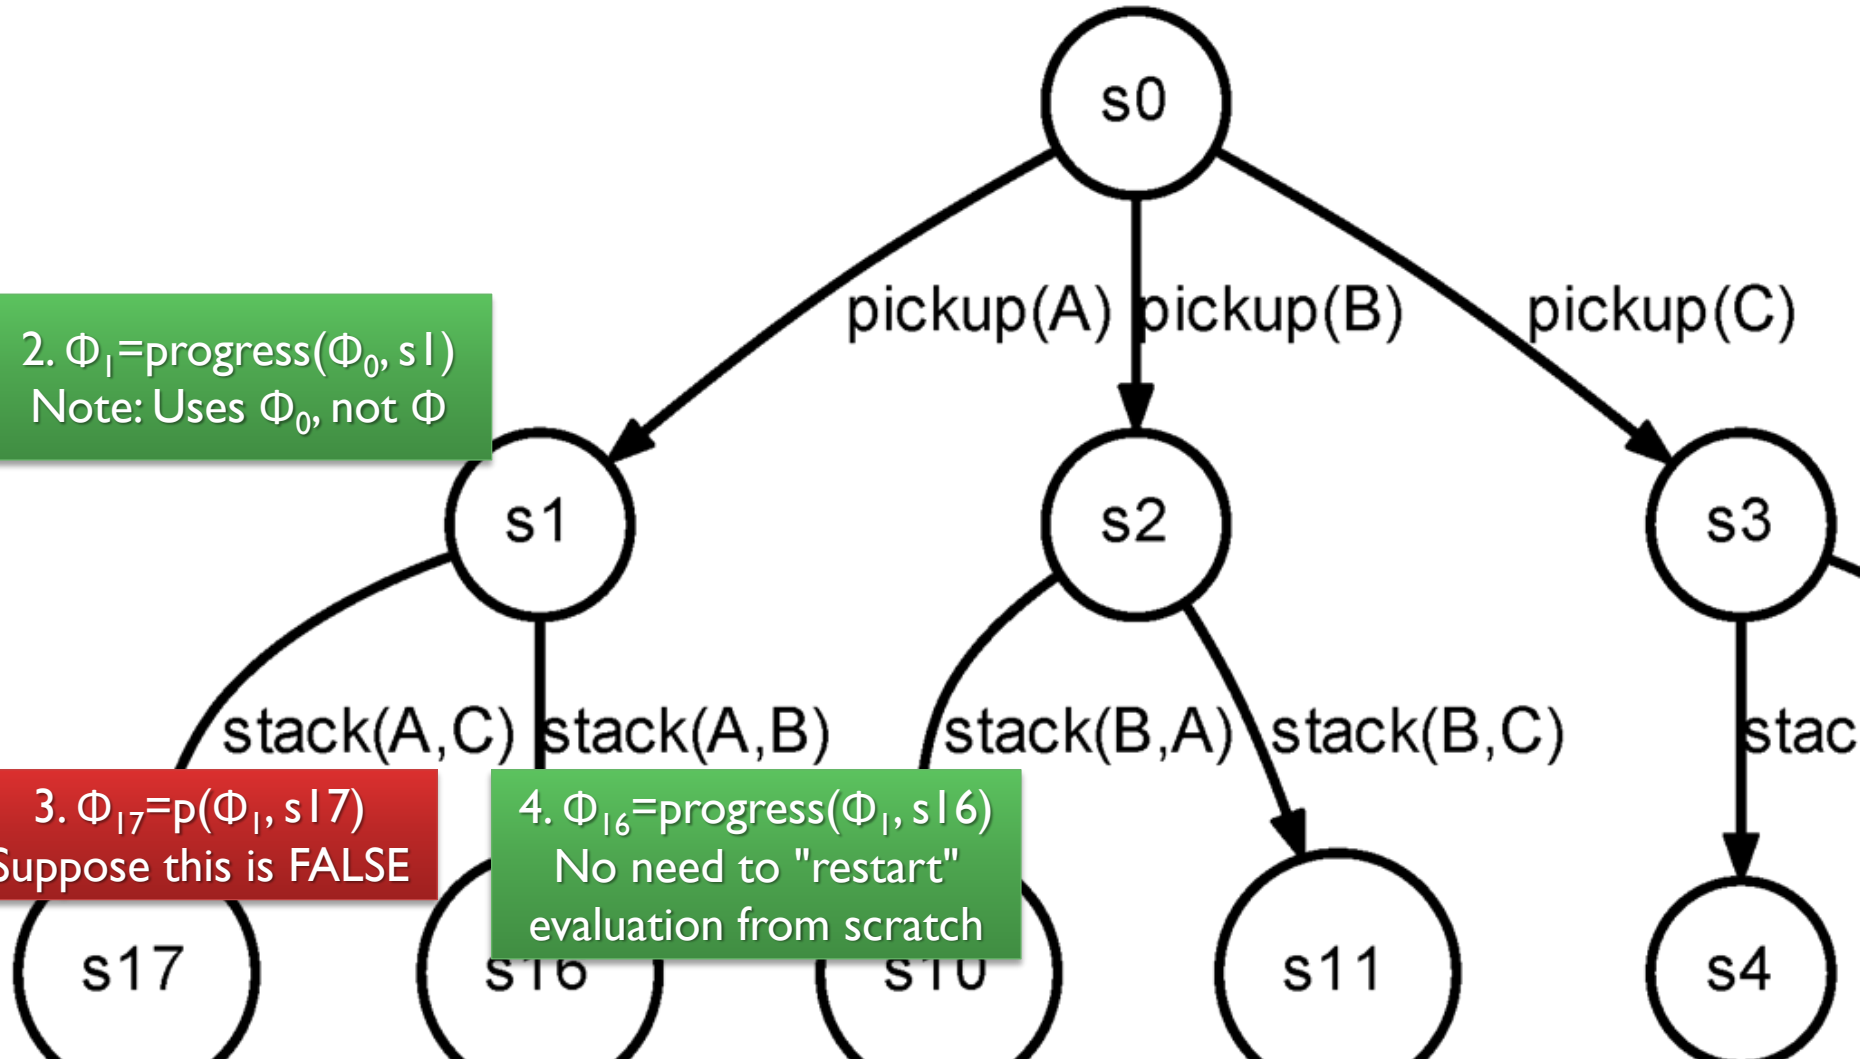
Φ = original control formula

1. Calculate and store $\Phi_0 = \text{progress}(\Phi, s_0)$
Resolves all references to values in state s_0 !

2. $\Phi_1 = \text{progress}(\Phi_0, s_1)$
Note: Uses Φ_0 , not Φ

3. $\Phi_{17} = p(\Phi_1, s_{17})$
Suppose this is FALSE

4. $\Phi_{16} = \text{progress}(\Phi_1, s_{16})$
No need to "restart"
evaluation from scratch



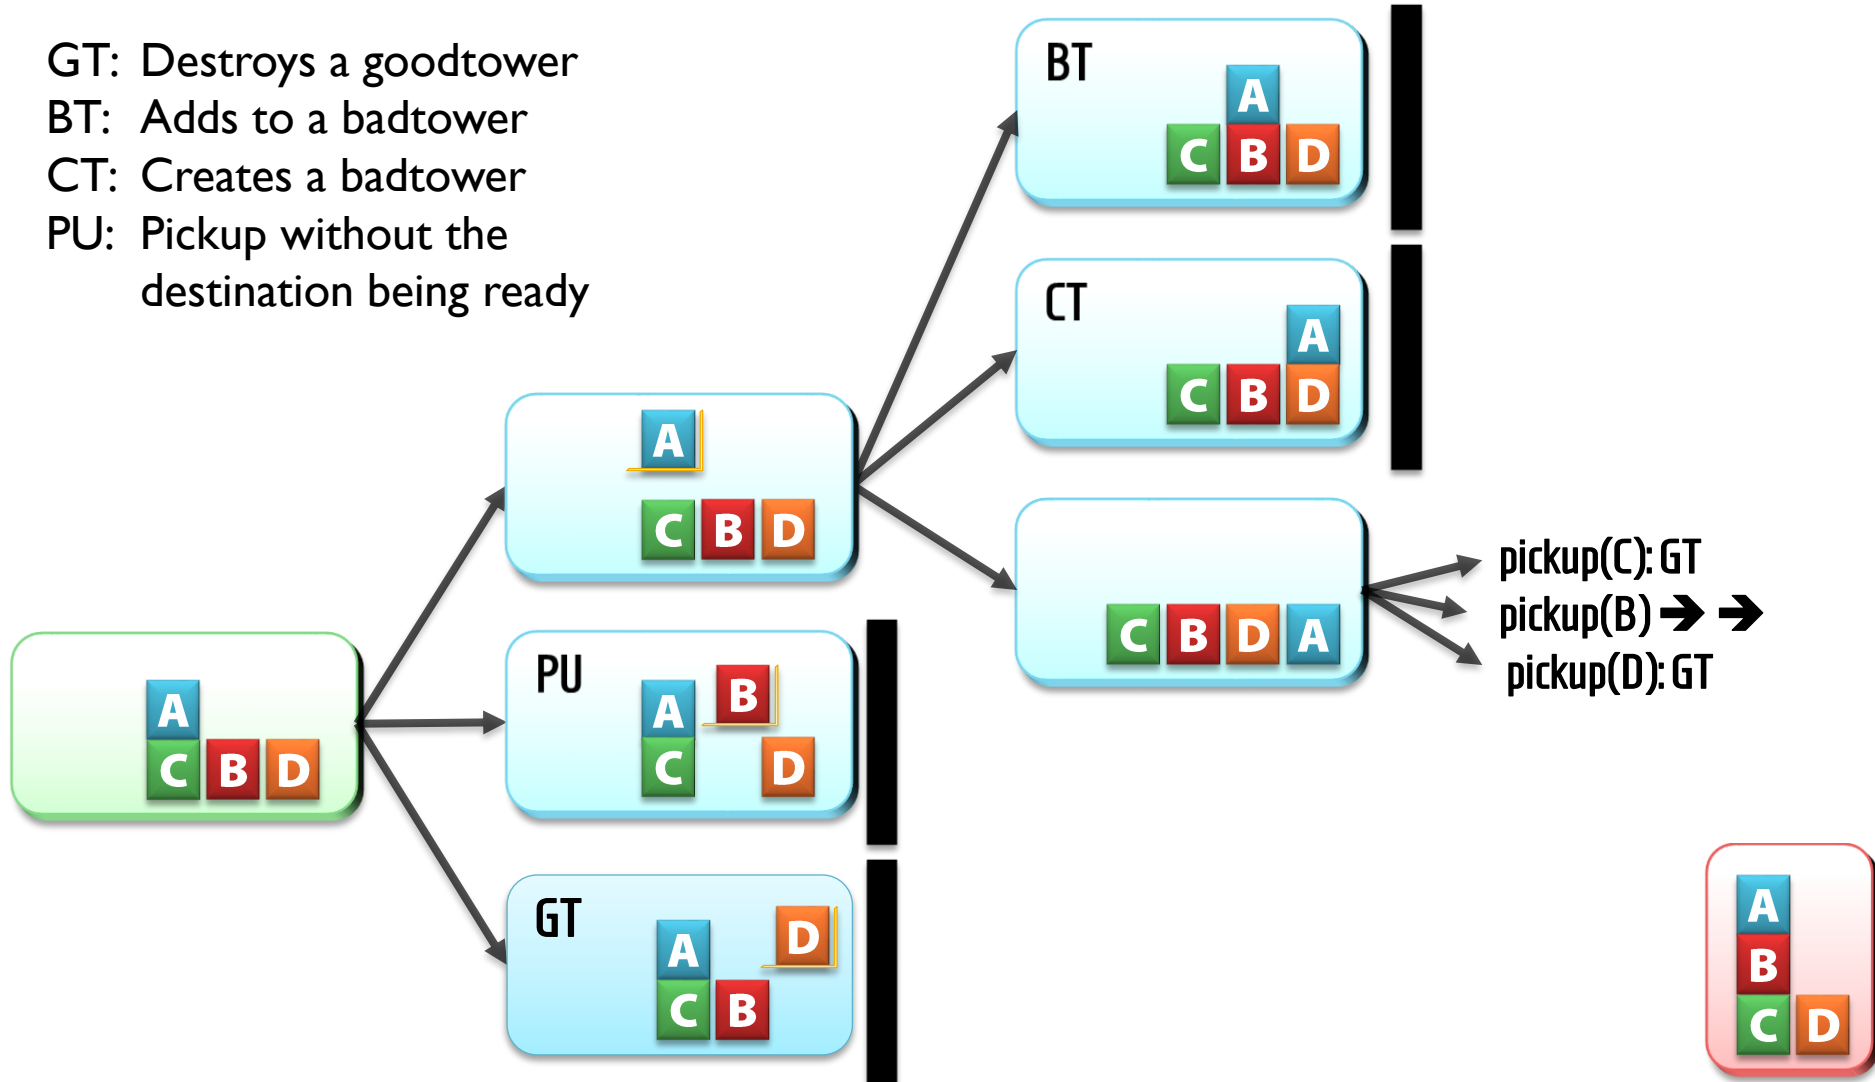
DFS with Pruning

GT: Destroys a goodtower

BT: Adds to a badtower

CT: Creates a badtower

PU: Pickup without the
destination being ready



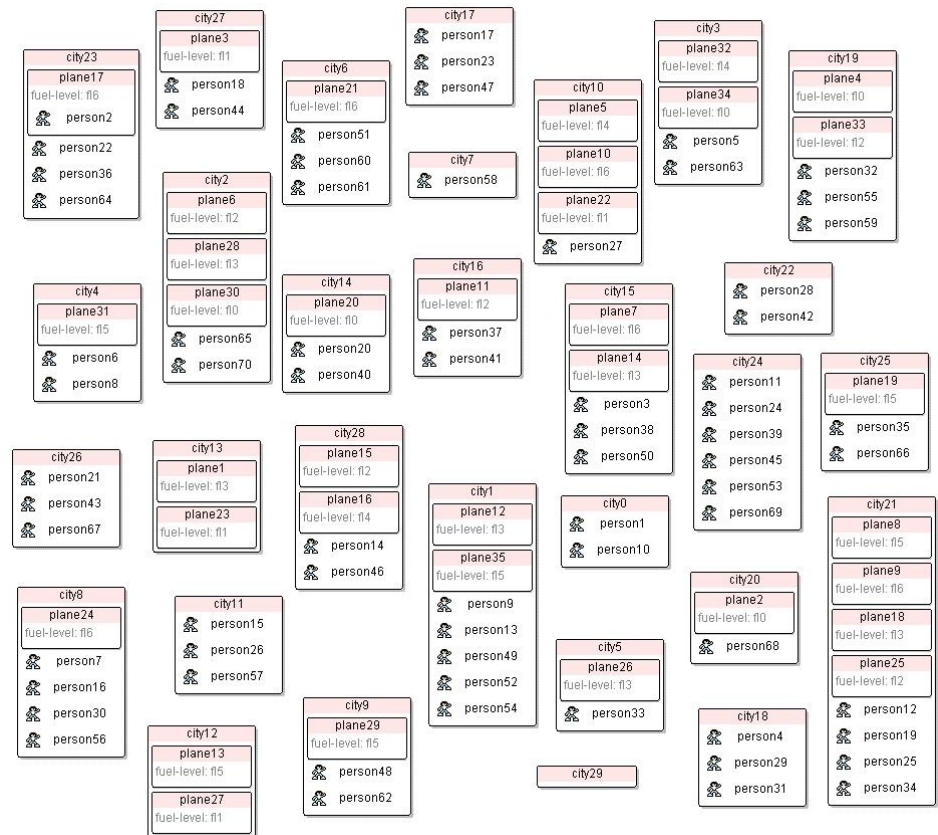
- 2000 International Planning Competition
 - **TALplanner** received the top award for a “hand-tailored” (i.e., domain-configurable) planner
- 2002 International Planning Competition
 - **TLplan** won the same award
- Both of them (as well as SHOP, an HTN planner):
 - Ran **several orders of magnitude** faster than the “fully automated” (i.e., not domain-configurable) planners
 - especially on large problems
 - Solved problems on which other planners ran out of time/memory
 - Required a **considerably greater modeling effort** for each planning domain

TALplanner: A demonstration

TALplanner Example Domain


■ Example Domain: ZenoTravel

- Planes move people between cities (board, debark, fly)
- Planes have limited fuel level; must refuel
- Example instance:
 - 70 people
 - 35 planes
 - 30 cities




ZenoTravel Problem Instance

- A smaller problem instance



What Just Happened?

- No additional domain knowledge specified yet!
 - Pure depth first...






- **First problem** in the example:
 - Passengers debark whenever possible.
 - Rule: "At any timepoint, if a passenger debarks, he is at his goal."
- **#control**:name "only-debark-when-in-goal-city"
 forall $t, person, aircraft$ [
 $[t]$ **in**($person, aircraft$) \rightarrow
 $[t+1]$ **in**($person, aircraft$) \vee
 exists $city$ [
 $[t]$ **at**($aircraft, city$) \wedge
 goal(**at**($person, city$))]]

$[t]$: "now"
 $[t+1]$: "next"

- **Second problem** in the example:

- Passengers board planes, even at their destinations
- Rule: "At any timepoint, if a passenger boards a plane, he was not at his destination."
- #**control**:name "only-board-when-necessary"
 forall $t, person, aircraft$ [
 ($[t]$ **!in**($person, aircraft$) \wedge
 $[t+1]$ **in**($person, aircraft$)) \rightarrow
 exists $city1, city2$ [
 $[t]$ **at**($person, city1$) \wedge
 goal(**at**($person, city2$)) \wedge
 $city1 \neq city2$]]

Zeno Travel, second attempt





What's Wrong This Time?



- Only constrained passengers
- Forgot to constrain airplanes
 - Which cities are reasonable destinations?
 - 1. A passenger's destination
 - 2. A place where a person wants to leave
 - 3. The airplane's destination

- #control :name "planes-always-fly-to-goal"
 forall t, aircraft, city [
 [t] at(aircraft, city) →
 ([t+1] at(aircraft, city)) |
 exists city2 [
 city2 != city &
 ([t+1] at(aircraft, city2)) &
 [t] reasonable-destination(aircraft, city2)]]
- #**define** [t] reasonable-destination(aircraft, city):
 [t] has-passenger-for(aircraft, city) |
 exists person [
 [t] at(person, city) &
 [t] in-wrong-city(person)] |
 goal(at(aircraft, city)) &
 [t] empty(aircraft) &
 [t] all-persons-at-their-destinations-or-in-planes]

Zeno Travel, third attempt





Conclusion

- Control Rules or Hierarchical Task Networks?
 - Both can be very efficient and expressive
 - If you have "recipes" for everything, HTN can be more convenient
 - Can be modeled with control rules, but not intended for this purpose
 - You have to forbid everything that is "outside" the recipe
 - If you have knowledge about "some things that shouldn't be done":
 - With control rules, the default is to "try everything"
 - Can more easily express localized knowledge about what should and shouldn't be done
 - Doesn't require knowledge of all the ways in which the goal can be reached