

CSC2542

Planning-Graph Techniques

Sheila McIlraith
Department of Computer Science
University of Toronto
Winter 2009

1

Administrative Announcements

- Tutorial tomorrow: Fri Jan 30 in BA2139, 1-2 PM
We will predominantly discuss the *course project*, but will also talk about possible reading topics, time permitting.
- If you haven't already met w/ me concerning your project, please make sure to email me to arrange a meeting time.
- Suggested readings for next week:
 - Read Part III and Chapter 9
 - Read a paper I will post shortly (no critique required).

2

END of Administrative Announcements

3

Acknowledgements

A number of the slides used in this course are modifications of Dana Nau's lecture slides for the textbook *Automated Planning*, licensed under the Creative Commons Attribution-NonCommercial-ShareAlike License: <http://creativecommons.org/licenses/by-nc-sa/2.0/>

Other slides are modifications of slides developed by Malte Helmert, Bernhard Nebel, and Jussi Rintanen.

I have also used some material prepared by Dan Weld, P@trick Haslum and Rao Kambhampati.

I would like to gratefully acknowledge the contributions of these researchers, and thank them for generously permitting me to use aspects of their presentation material.

4

History

- **GraphPlan** (Blum & Furst, 1995) was the first planner to use **planning-graph techniques**
- Before GraphPlan came out, most planning researchers were working on partial order planners (plan space planners *)
 - POP, SNLP, UCPOP, etc.
- GraphPlan caused a sensation because it was so much faster
- Many subsequent planning systems have used ideas from it either directly as close descendants of GraphPlan or by using the Planning Graph representation in some guise to improve the encoding of the planning problem most notably for SAT-based planning.

* Sometimes referred to as "PSP" planners, but "PSP" used for "partial satisfaction planners", nowadays

6

History

But most importantly...

- GraphPlan's place in history is secured by its critical role in the development of *reachability heuristics** for heuristic search planners by approximating the search tree rooted at a given state
- Heuristic search planners have dominated the "satisficing planner" track of IPC planning competitions for the last 8 years.

* *Reachability heuristics aim to estimate the cost of a plan between the current search state and a goal state. We will talk about these more in the weeks to come.*

7

Outline

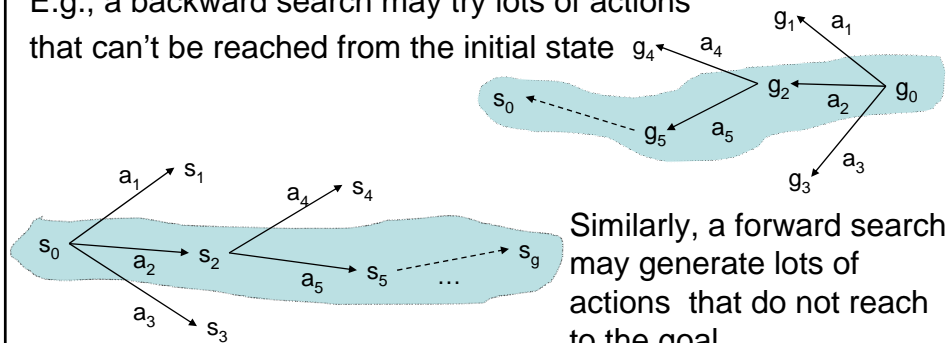
- Motivation
- The Graphplan algorithm
- Planning graphs
 - example
- Mutual exclusion
 - example (continued)
- Solution extraction
 - example (continued)
- Discussion

8

Motivation

- A big source of inefficiency in search algorithms is the *branching factor*
 - the number of children of each node

E.g., a backward search may try lots of actions that can't be reached from the initial state



9

One way to reduce branching factor

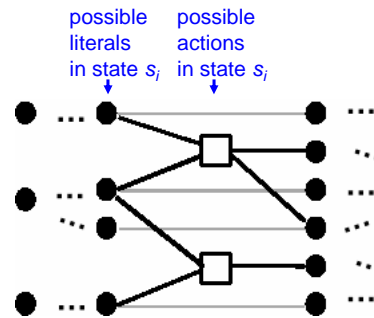
- First create a *relaxed problem*
 - Remove some restrictions of the original problem
 - Want the relaxed problem to be easy to solve (polynomial time) ← **IMPORTANT**
 - The solutions to the relaxed problem will include all solutions to the original problem
- Then do a modified version of the original search
 - Restrict its search space to include only those actions that occur in solutions to the relaxed problem

10

Graphplan

procedure Graphplan:

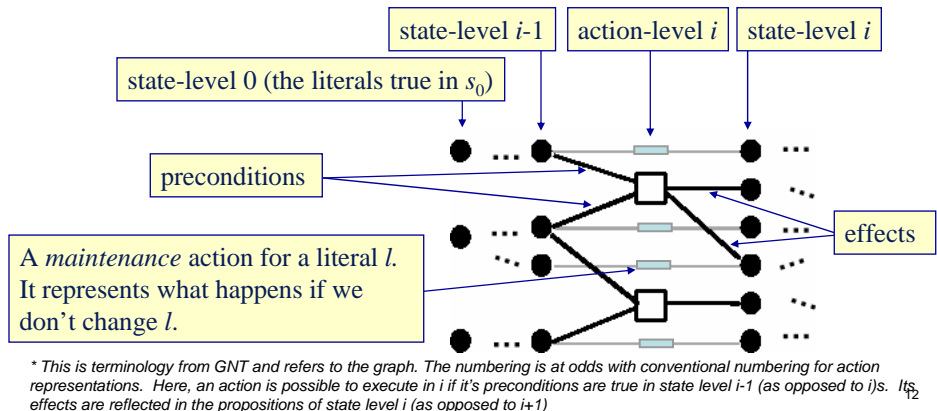
- for $k = 0, 1, 2, \dots$
 - *Graph Expansion:*
 - create a “planning graph” that contains k “levels” relaxed problem
 - Check whether the planning graph satisfies a necessary (but insufficient) condition for plan existence
 - If it does, then
 - do *Solution Extraction:*
 - backward search, modified to consider only the actions in the planning graph
 - if we find a solution, then return it



11

The Planning Graph

- Search space for a relaxed version of the planning problem
- Alternating layers of ground literals and actions
 - Nodes at action-level i : actions that might be possible to execute at time i^*
 - Nodes at state-level i : literals that might possibly be true at time i
 - Edges: preconditions and effects



Example

- Due to Dan Weld, Univ. Washington [Weld, AIM-09]
- Suppose you want to prepare dinner as a surprise for your sweetheart (who is asleep)

$s_0 = \{\text{garbage, cleanHands, quiet}\}$

$g = \{\text{dinner, present, } \neg\text{garbage}\}$

Action	Preconditions	Effects
cook()	cleanHands	dinner
wrap()	quiet	present
carry()	none	\neg garbage, \neg cleanHands
dolly()	none	\neg garbage, \neg quiet

Also have the maintenance actions: one for each literal

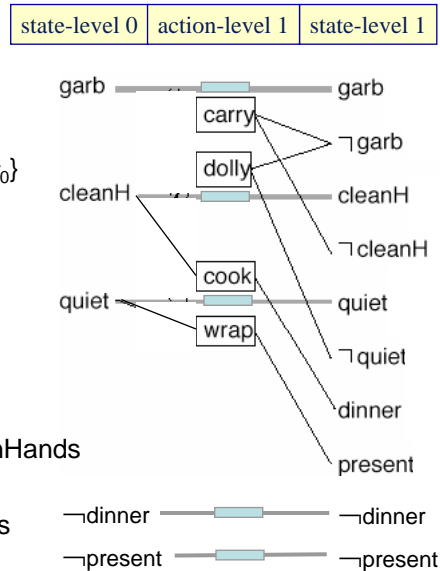
Example (continued)

- state-level 0:
 {all atoms in s_0 } \cup
 {negations of all atoms not in s_0 }
- action-level 1:
 {all actions whose preconditions
 are satisfied and non-mutex in s_0 }
- state-level 1:
 {all effects of all of the
 actions in action-level 1}

Action Preconditions Effects

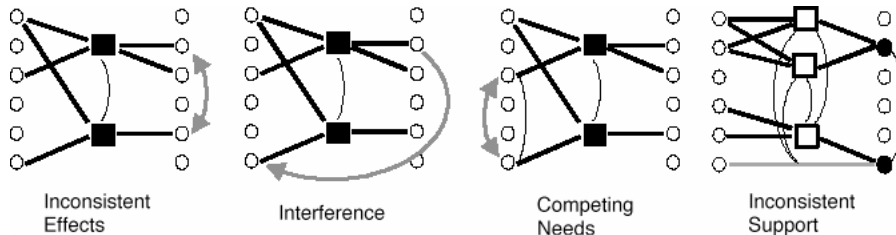
cook()	cleanHands	dinner
wrap()	quiet	present
carry()	none	\neg garbage, \neg cleanHands
dolly()	none	\neg garbage, \neg quiet

Also have the maintenance actions



14

Mutual Exclusion



- Two actions at the same action-level are mutex if
 - Inconsistent effects*: an effect of one negates an effect of the other
 - Interference*: one deletes a precondition of the other
 - Competing needs*: **they have mutually exclusive preconditions**
- Otherwise they don't interfere with each other
 - Both may appear in a solution plan
- Two literals at the same state-level are mutex if
 - Inconsistent support*: one is the negation of the other, **or all ways of achieving them are pairwise mutex**

Recursive propagation of mutexes

15

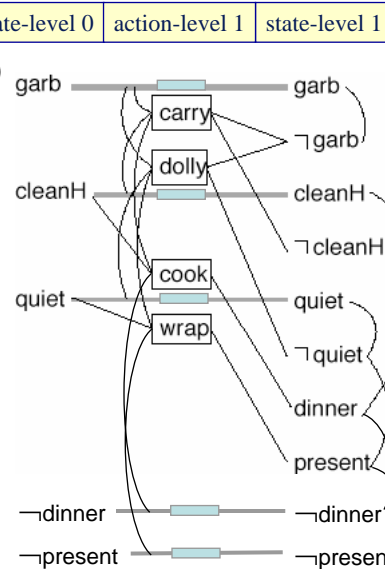
Example (continued)

- Augment the graph to indicate mutex
- *carry* is mutex with the maintenance action for *garbage* (inconsistent effects)
- *dolly* is mutex with *wrap*
 - interference
- \neg *quiet* is mutex with *present*
 - inconsistent support
- each of *cook* and *wrap* is mutex with a maintenance operation

Action Preconditions Effects

cook()	cleanHands	dinner
wrap()	quiet	present
carry()	<i>none</i>	\neg garbage, \neg cleanHands
dolly()	<i>none</i>	\neg garbage, \neg quiet

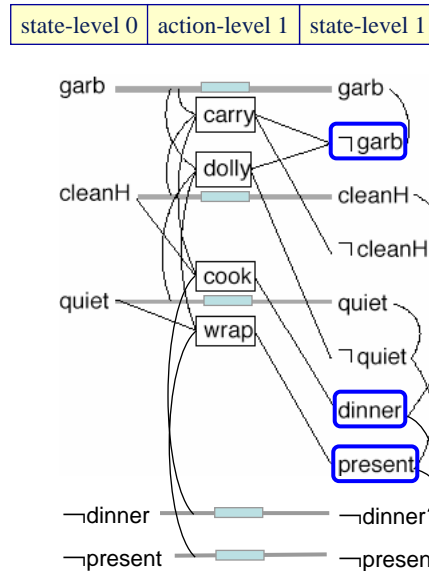
Also have the maintenance actions



16

Example (continued)

- Check to see whether there's a possible solution
- Recall that the goal is
 - $\{\neg$ *garbage*, *dinner*, *present* $\}$
- Note that in state-level 1,
 - All of them are there
 - None are mutex with each other
- Thus, there's a chance that a plan exists
- Try to find it
 - Solution extraction



17

Recall what the algorithm does

procedure Graphplan:

- for $k = 0, 1, 2, \dots$
 - *Graph Expansion*:
 - create a “planning graph” that contains k “levels”
 - Check whether the planning graph satisfies a necessary (but insufficient) condition for plan existence
 - If it does, then
 - do *Solution Extraction*:
 - backward search, modified to consider only the actions in the planning graph
 - if we find a solution, then return it

18

Solution Extraction

The set of goals we are trying to achieve

The level of the state s_i

procedure Solution-extraction(g, i)

if $i=0$ then return the solution

for each literal l in g

nondeterministically choose an action

to use in state s_{i-1} to achieve l

if any pair of chosen actions are mutex

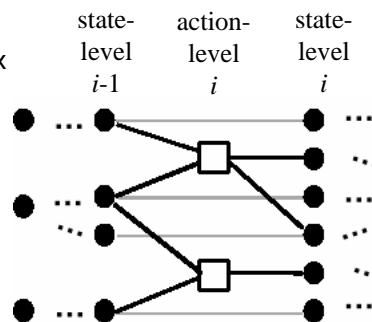
then backtrack

$g' := \{\text{the preconditions of the chosen actions}\}$

Solution-extraction($g', i-1$)

end Solution-extraction

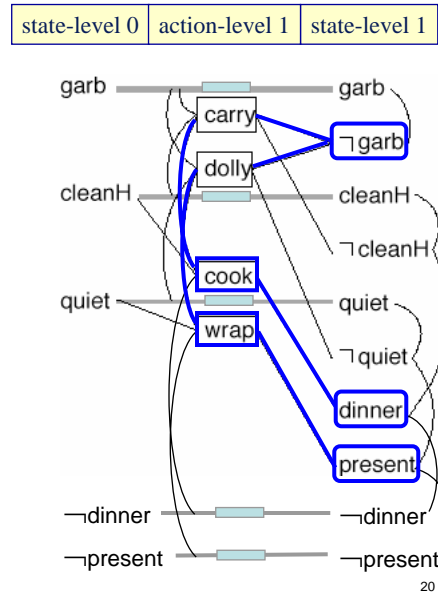
A real action or a maintenance action



19

Example (continued)

- Two sets of actions for the goals at state-level 1
- Neither of them works
 - Both sets contain actions that are mutex



20

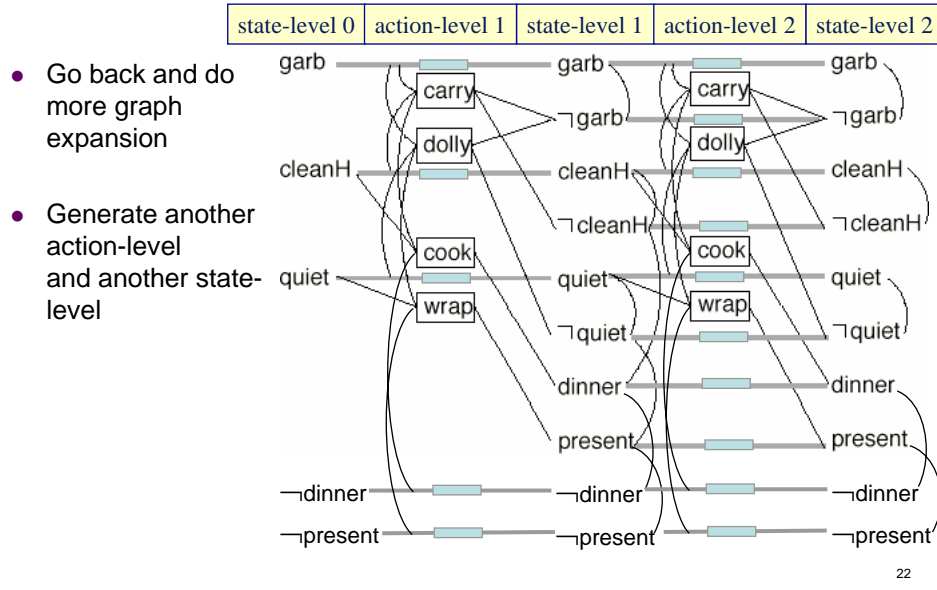
Recall what the algorithm does

procedure Graphplan:

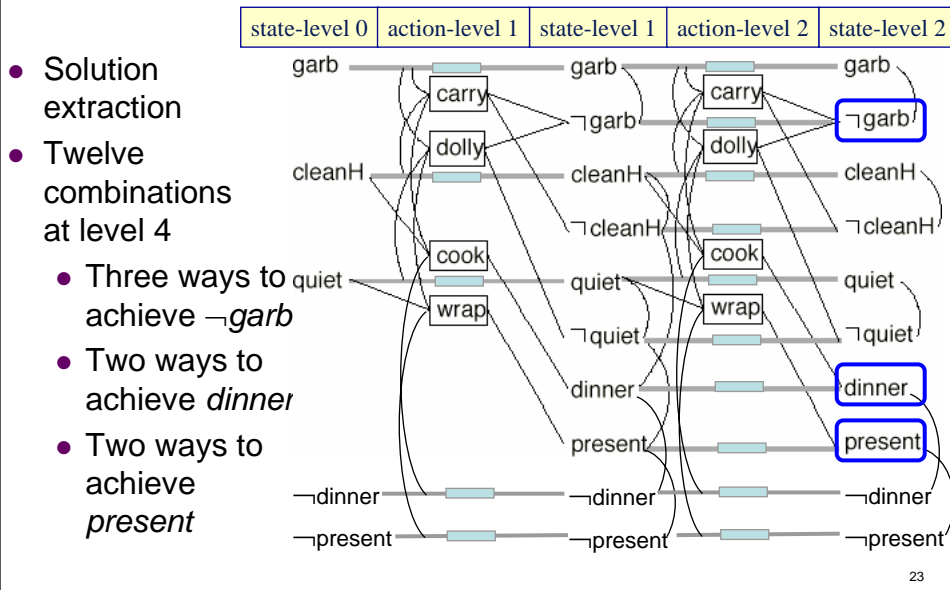
- for $k = 0, 1, 2, \dots$
 - *Graph Expansion*:
 - create a “planning graph” that contains k “levels”
 - Check whether the planning graph satisfies a necessary (but insufficient) condition for plan existence
 - If it does, then
 - do *Solution Extraction*:
 - backward search, modified to consider only the actions in the planning graph
 - if we find a solution, then return it

21

Example (continued)

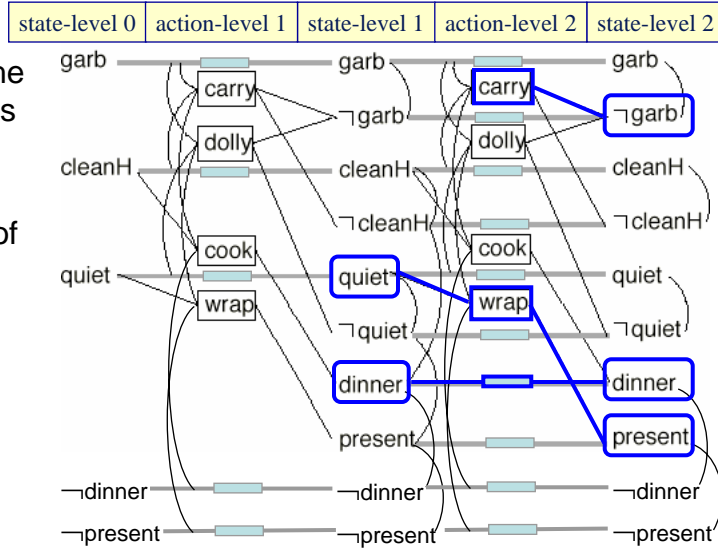


Example (continued)



Example (continued)

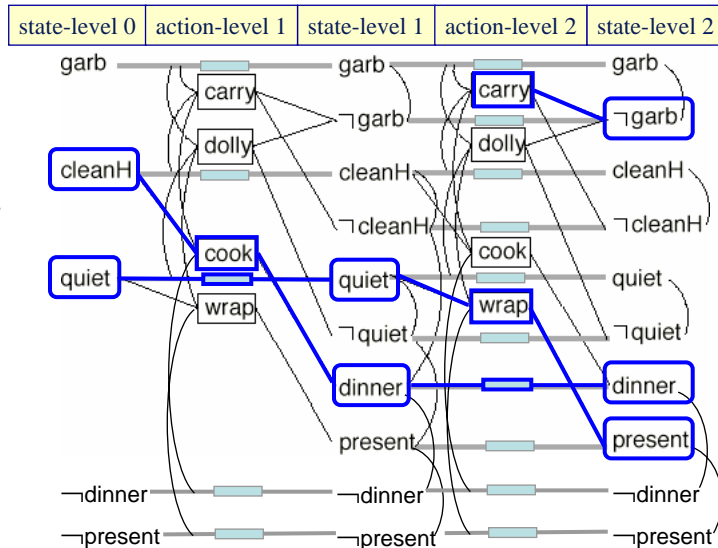
- Several of the combinations look OK at level 2
- Here's one of them



24

Example (continued)

- Call Solution-Extraction recursively at level 2
- It succeeds
- Solution whose *parallel length* is 2



25

Observation

- The solution is a sequence of *sets of actions* (as opposed to simply a sequence of actions)
- To generate a sequential plan, the solution can be serialized (in a variety of ways)

26

Properties of GraphPlan

- Graphplan is sound and complete
 - If Graphplan returns a plan, then it is a solution to the planning problem.
 - If solutions exist, then Graphplan will return one of them.
- The size of the planning graph Graphplan generates is polynomial in the size of the planning problems.
- Solutions extraction is still exponential in the worst case.
For many problems, Graph Expansion dominates problem solving time.
- The planning algorithm always terminates
 - There is a fixpoint on the number of levels of the planning graphs such that the algorithm either generates a solution or returns failure

27

Further Analysis

- (+) The backward-search part of Graphplan—which is the hard part—will only look at the actions in the planning graph – a smaller search space than the original one.
- (-) To generate the planning graph, Graphplan creates a huge number of ground atoms
- Many of them may be irrelevant
- Can alleviate (but not eliminate) this problem by assigning data types to the variables and constants

28

Optimizations and Extensions

Optimizations to the original Graphplan model:

- Improvements to solution extraction (e.g., forward checking, memoization, EBL)
- Improvements to graph expansion (e.g., closed-world assumption, compilation of action schemata w/ type analysis, in-place graph expansion)

Most of the **extensions** relate to language extensions:

- E.g., Universal quantification, conditional effects, negated preconditions and goals, ...
- Addressed by compilation → large increases in problem size, or
- Addressed by changing (complicating) the expansion, mutex determination and extraction

See [Weld, AIM-99] posted on our web page.

29

Optimizations and Extensions

Optimizations to the original Graphplan model:

- Improvements to solution extraction (e.g., forward checking, memoization, EBL)
- Improvements to graph expansion (e.g., closed-world assumption, compilation of action schemata w/ type analysis, in-place graph expansion)

Most of the **extensions** relate to language extensions:

- E.g., Universal quantification, conditional effects, negated preconditions and goals, ...
- Addressed by compilation → large increases in problem size, or
- Addressed by changing (complicating) the expansion, mutex determination and extraction

See [Weld, AIM-99] posted on our web page.

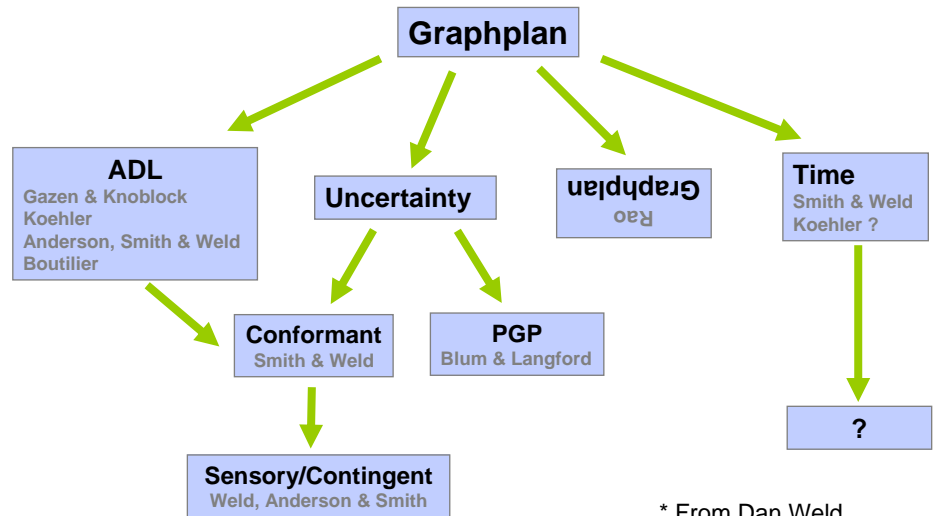
31

Graphplan and Its Descendants

- **Graphplan** (Blum & Furst, 1995) – C implementation
- **IPP** (Koehler et al. 1997) – highly optimized C implementation, extended to handle universal quantification and conditional effects
- **STAN** (Long & Fox, 1998)- highly optimized C, implementation – uses in-place graph rep'n and performs sophisticated type analysis
- **SGP** (Weld et al., 1998) – Lisp implementation – extended to handle universal quantification, conditional effects, and uncertainty

32

Perverting Graphplan*



* From Dan Weld