Partial Satisfaction Planning

J. Benton, Romeo Sanchez
Minh Do (PARC)
Subbarao Kambhampati

ARIZONA STATE UNIVERSITY

[AAAI 04; ICAPS 05; IJCAI 05]

Yochan Research Group

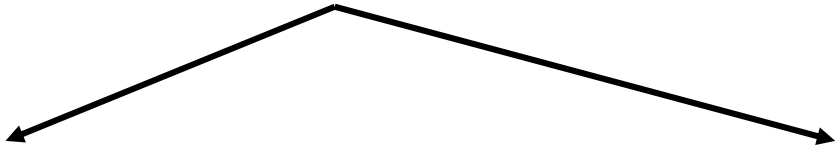

Plan-Yochan

- Automated Planning
  - Foundational work in classical planning
  - Heuristics for scaling up a wide spectrum of plan synthesis problems
  - Applications to manufacturing, biological pathway discovery, web services, autonomic computing

Recent Alumni

B. Srivastava; IBM IRL
M. Do; Xerox PARC
T. Zimmerman; CMU RI
R. Sanchez; USC/ISI

Db-Yochan

- Information Integration
  - Mediator frameworks that are adaptive to the sources and users.
  - Applications to Bio-informatics, Archaelogical informatics

Z. Nie; Microsoft Research
T. Hernandez; Amazon
U. Nambiar; SDSC/UC Davis

In many real world planning tasks, the agent often has more goals than it has resources to accomplish.

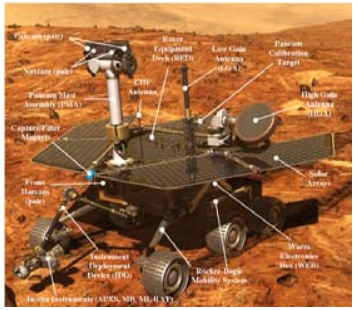

Currently *humans* are forced to pick goal subsets

Example: Rover Mission Planning (MER)

Military logistics

Most replanning problems (*)

Need automated support for Over-subscription/Partial Satisfaction Planning

Actions have execution costs, goals have utilities, and the objective is to find the plan that has the highest net benefit.

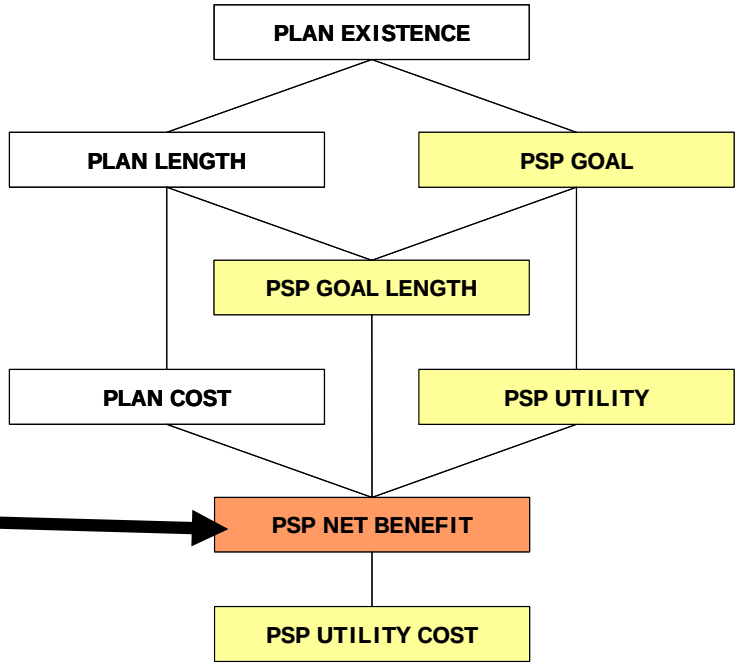

A spectrum of approaches for PSP-Net Benifit

Optimal Approaches

- **Deterministic MDPs**
  - Reward of a state is equal to the utility of the goals that hold in it.
    - Need to avoid collecting rewards for a goal more than once
  - Optimal, but *SLOW*
- **Optiplan**
  - Integer programming based STRIPS planner
    - Optimal for a given plan length

Heuristic Approaches

- **AltAlt^{ps} / AltAwlt^{ps}**
  - Selects the "objectives" up front heuristically
    - Uses a clever modification of relaxed plan heuristic
    - Not optimal, but fast
- **Sapa^{ps} / Sapa^{Mps}**
  - Models PSP as heuristic search. Can be optimal given admissible heuristics.
  - Sapa^{Mps} can handle numerical goals and degrees of satisfaction

Adapting PG heuristics for PSP

- **Challenges:**
  - **Need to propagate costs on the planning graph**
  - The exact set of goals are not clear
    - Interactions between goals
    - Obvious approach of considering all 2^n goal subsets is *infeasible*

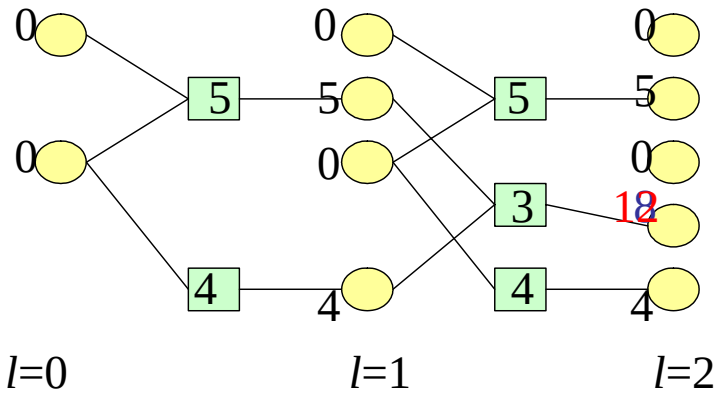

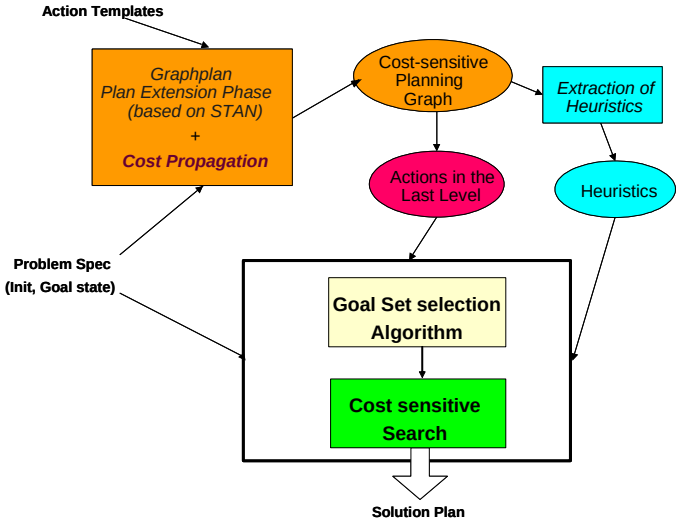

- **Idea:** Select a subset of the top level goals upfront
- Challenge: Goal interactions

- Approach: Estimate the net benefit of each goal in terms of its utility minus the cost of its relaxed plan
  - Bias the relaxed plan extraction to (re)use the actions already chosen for other goals

Some Empirical Results for AltAlt^{ps}

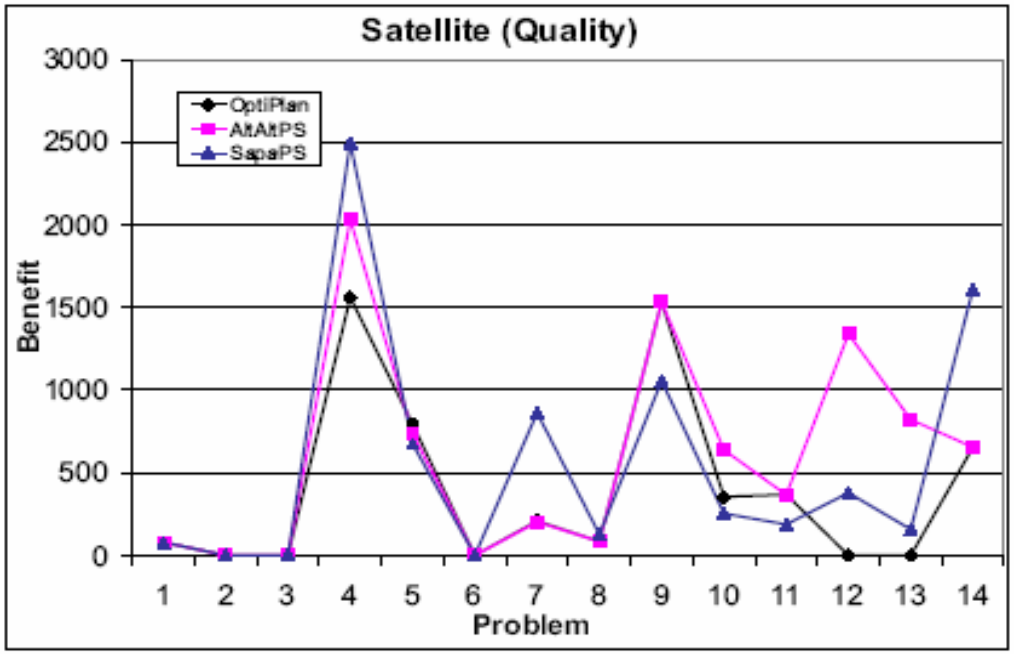

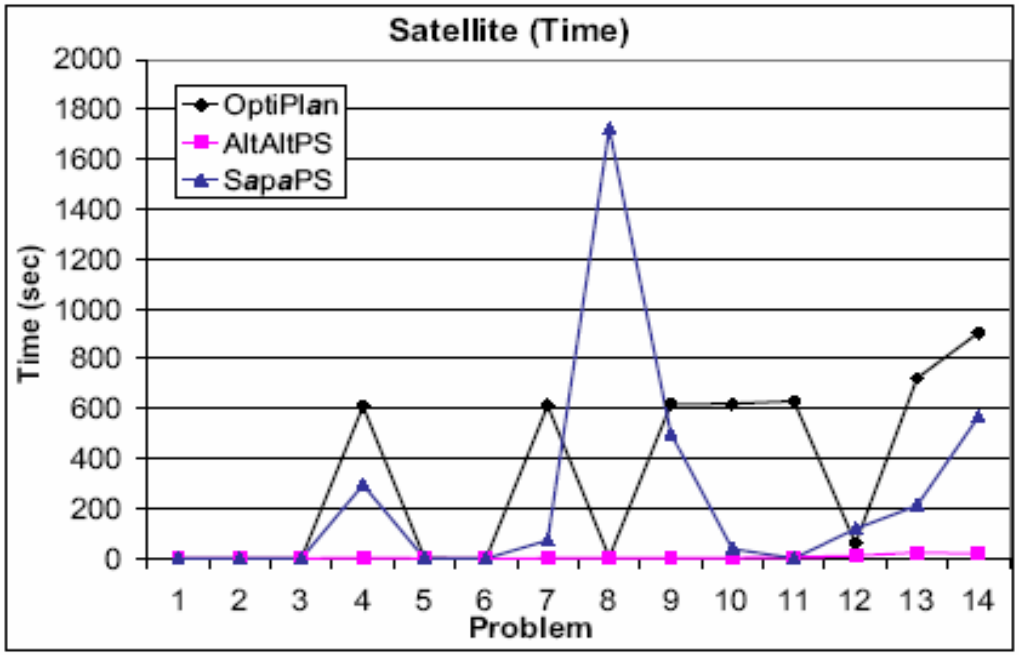

Exact algorithms based on MDPs don't scale at all

[AAAI 2004]

Overcoming Complex Interactions: AltAwlt [ICAPS-05]

- **Problems with Goal Selection Procedure**
  1. Ignores group interactions
  2. Ignores negative interactions

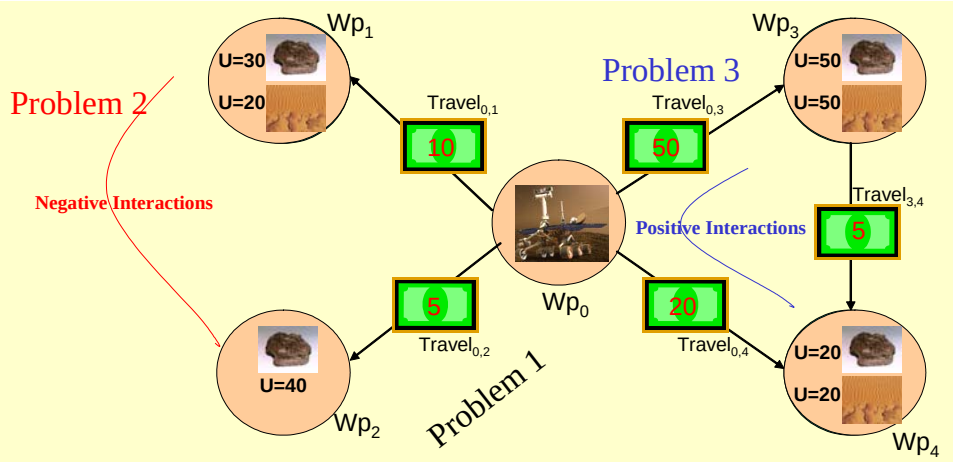

- **Ideas:**
  1. Consider multiple groups of sub-goals during the selection process
  2. Add penalty costs for ignoring negative interactions based on mutex analysis

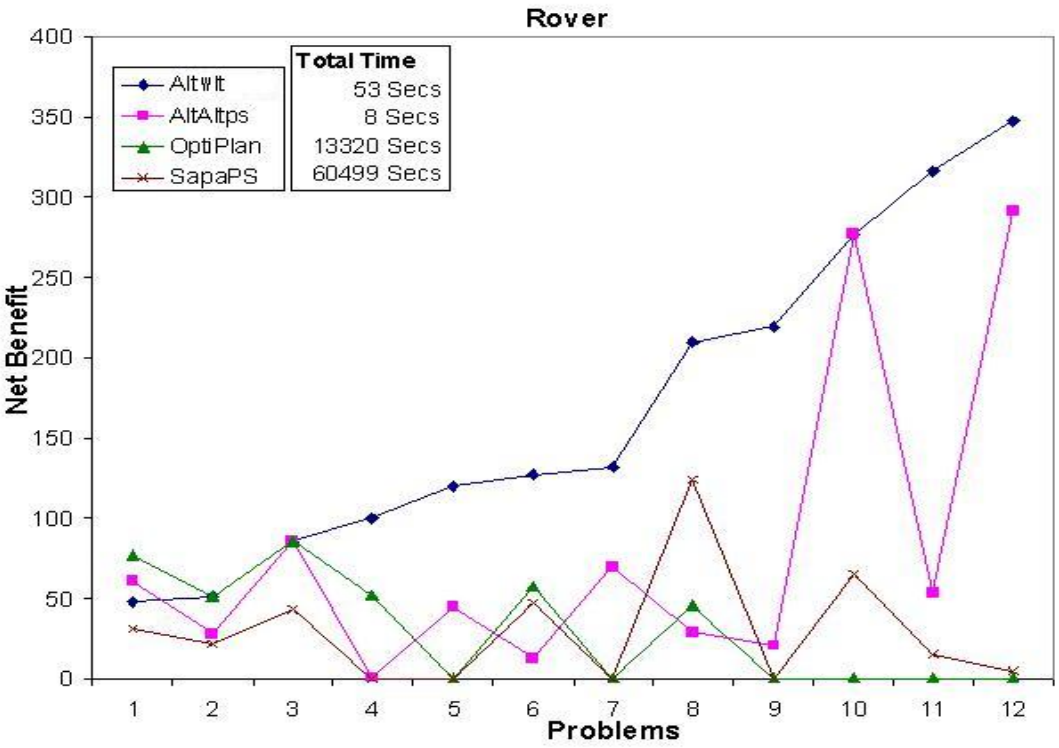

PSP+MTP=SAPA$^{\mathcal{M}ps}$

[IJCAI 2005]

- In MTP, PSP will involve
  - Partial Degree of satisfaction
    - If you can't give me 1000$, give me half at least
    - Need to track costs for various intervals of a numeric quantity ☹
  - Delayed Satisfaction
    - If you submit the homework past the deadline, you will get penalty points

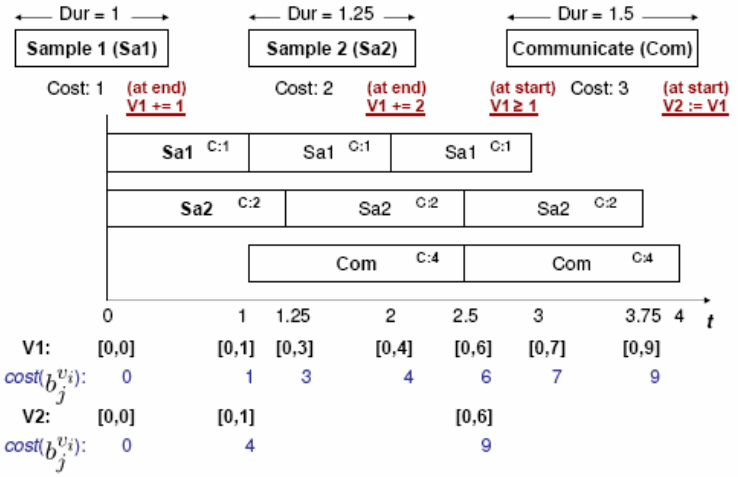

Figure 3: The RTPG for our example. Our actions are defined above it.

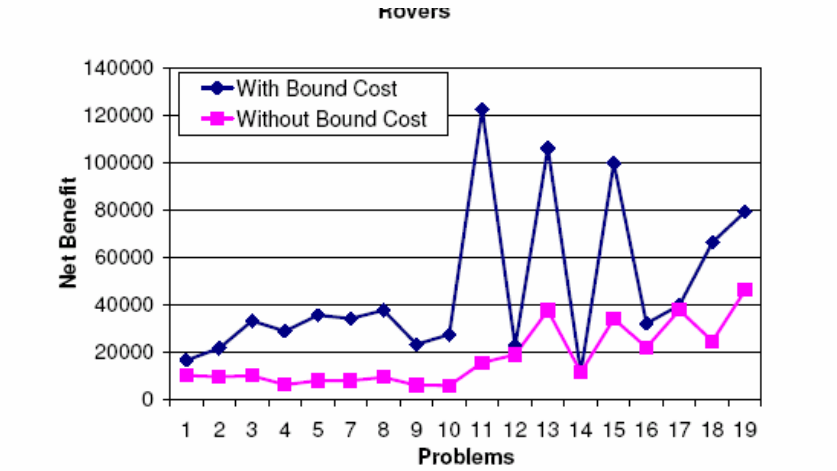

Figure 4: Comparison of utilities for our rovers domain