

Planning through Stochastic Local Search and Temporal Action Graphs in LPG

By Alfonso Gervini, Alessandro Saetti, Ivan Serina

Presented by Evan Damron

Overview

- ▶ LPG is a planning graph-based planner, capable of solving plans where actions take time (temporal), and consume resources (numerical).
- ▶ LPG locally searches the space of linear action graphs by adapting the SAT solver called Walk-SAT to the problem, calling it Walkplan.
- ▶ LPG was the best planner in the 3rd International Planning Competition (IPC-3) in 2002.

Action Graphs

- ▶ Action graphs are subgraphs of the planning graph, representing partially ordered plans.
- ▶ Action graphs contain a dummy start and end step.
- ▶ If an action is in an action graph, then so are its preconditions and effects.
- ▶ Action graphs can contain inconsistencies:
 - ▶ Fact nodes not supported by an action.
 - ▶ Mutexed actions.
- ▶ A solution graph is an action graph without inconsistencies.

Persistent Mutexes

- ▶ Mutexes decrease monotonically as the levels of the planning graph increase. i.e. all mutexes are guaranteed to be in all previous levels if the facts/actions involved in that mutex are in these levels.
- ▶ The *fixed-point* level of the graph is the point where nodes and mutexes will remain the same at all further levels.
- ▶ Mutexes at the fixed-point level are called *persistent mutexes*.
- ▶ In Graphplan, mutexes are computed a level at a time, which can be computationally very expensive.
- ▶ LPG speeds this up by only computing persistent mutexes.

Calculating Persistent Mutexes

ComputeMutexFacts($I, \mathcal{O}$)

Input: An initial state (I) and all ground operator instances ($\mathcal{O}$);

Output: A set of persistent mutex relations between facts (M).

```
1   $F^* \leftarrow I; F \leftarrow \emptyset;$ 
2   $M \leftarrow \emptyset; M^* \leftarrow \emptyset; A \leftarrow \emptyset;$ 
3  while  $F^* \neq F \vee M^* \neq M$ 
4     $F \leftarrow F^*; M \leftarrow M^*;$ 
5    forall  $a \in \mathcal{O}$  such that  $Pre(a) \subseteq F^*$  and  $\neg(\exists p, q \in Pre(a) \wedge (p, q) \in M^*)$ 
6       $New(a) \leftarrow Add(a) - F^*;$ 
7      forall  $f \in New(a)$ 
8        forall  $h \in Del(a)$ 
9           $M^* \leftarrow M^* \cup \{(f, h), (h, f)\};$  /* Potential mutex relation */
10       forall  $(p, q) \in M^*$  such that  $p \in Pre(a)$  and  $q \notin Del(a)$ 
11          $M^* \leftarrow M^* \cup \{(f, q), (q, f)\};$  /* Potential mutex relation */
12       if  $a \notin A$  then
13         forall  $p, q \in Add(a)$  such that  $(p, q) \in M^*$ 
14            $M^* \leftarrow M^* - \{(p, q), (q, p)\};$  /* Invalid mutex relation */
15        $L \leftarrow Add(a) - New(a);$ 
16       forall  $(i, q) \in M^*$  such that  $i \in L$ 
17         if  $q \notin Del(a) \wedge \neg(\exists p \in Pre(a) \wedge (p, q) \in M^*)$  then
18            $M^* \leftarrow M^* - \{(i, q), (q, i)\};$  /* Invalid mutex relation */
19        $F^* \leftarrow F^* \cup New(a);$ 
20        $A \leftarrow A \cup \{a\};$ 
21  return  $M$ .
```

Check if F1 is a positive effect for a and F2 is a negative effect for a .

Check if F1 is a positive effect for a and F2 is mutexed with a precondition of a .

F1 and F2 both belong to the add effects of an action.

F1 is an add effect of a , F2 is not deleted by a , and F2 isn't mutexed with any precondition of a .

Calculating Persistent Mutexes

ComputeMutexActions($M, \mathcal{O}$)

Input: A set of mutex relations between facts (M) and all ground operator instances ($\mathcal{O}$);

Output: A set of persistent mutex relations between actions (N).

1. $N \leftarrow \emptyset$; $\mathcal{O}^* \leftarrow \mathcal{O}$ extended with the no-op of every fact;
2. **forall** $(p, q) \in M$
3. **forall** $a \in \mathcal{O}^*$ such that $p \in \text{Pre}(a)$
4. **forall** $b \in \mathcal{O}^*$ such that $q \in \text{Pre}(b)$
5. $N \leftarrow N \cup \{(a, b), (b, a)\}$; /* Competing needs */
6. **forall** $a \in \mathcal{O}^*$
7. **forall** $p \in \text{Pre}(a)$
8. **forall** $b \in \mathcal{O}$ such that $p \in \text{Del}(b)$
9. $N \leftarrow N \cup \{(a, b), (b, a)\}$; /* Interference */
10. **forall** $p \in \text{Add}(a)$
11. **forall** $b \in \mathcal{O}$ such that $p \in \text{Del}(b)$
12. $N \leftarrow N \cup \{(a, b), (b, a)\}$; /* Inconsistent effects */
13. **return** N .

Check if any preconditions of two actions are mutexed.

Check if one action deletes a precondition of another.

Check if one action deletes an add effect of another.

Walkplan

- ▶ Walkplan first computes the planning graph and persistent mutexes.
- ▶ Walkplan then creates an initial action graph. The default contains only the no-ops of the facts in the initial state until the fixed-point level (with the start and end actions).
- ▶ At each step, Walkplan will choose an inconsistency and attempt to resolve it.
 - ▶ If the inconsistency is an unsupported fact:
 - ▶ Remove an action that needs this fact, OR
 - ▶ Add an action to produce this fact.
 - ▶ If the inconsistency is a mutex, remove one of the actions in the mutex relationship.
- ▶ Walkplan has a noise parameter p to prevent getting stuck in local optima.
- ▶ The *neighborhood* of an inconsistency in an action plan are all the possible action plans that can result from fixing the inconsistency.

Walkplan pseudocode


Walkplan($\Pi, max_steps, max_restarts, p$)

Input: A planning problem Π , the maximum number of search steps max_steps , the maximum number of search restarts $max_restarts$, a noise factor p ($0 \leq p \leq 1$).

Output: A solution graph representing a plan solving Π or **fail**.

1. **for** $i \leftarrow 1$ **to** $max_restarts$ **do**
2. $\mathcal{A} \leftarrow$ an initial A-graph derived from the planning graph of Π ;
3. **for** $j \leftarrow 1$ **to** max_steps **do**
4. **if** $\mathcal{A}$ is a solution graph **then**
5. **return** $\mathcal{A}$
6. $\sigma \leftarrow$ an inconsistency in $\mathcal{A}$;
7. $N(\sigma, \mathcal{A}) \leftarrow$ neighborhood of $\mathcal{A}$ for σ ;
8. **if** $\exists \mathcal{A}' \in N(\sigma, \mathcal{A})$ such that the quality of $\mathcal{A}'$ is not worse than the quality of $\mathcal{A}$
9. **then** $\mathcal{A} \leftarrow \mathcal{A}'$ (if there is more than one $\mathcal{A}'$ -graph, choose randomly one)
10. **else if** $random < p$ **then**
11. $\mathcal{A} \leftarrow$ an element of $N(\sigma, \mathcal{A})$ randomly chosen
12. **else** $\mathcal{A} \leftarrow$ best element in $N(\sigma, \mathcal{A})$;
13. **return fail**.

Walkplan Example




```
graph TD
    start([start]) --> read[read]
    read --> write[write]
    write --> read
    read --> read
    write --> write
```



Neighborhood: Add supporting action *eat(cake)* or no-op

Walkplan Example




Inconsistency count: 2

Selected Inconsistency: *eat(cake)* mutex with *-eaten(cake)* no-op

Neighborhood: Remove *eat(cake)* or *-eaten(cake)* no-op

Walkplan Example




Inconsistency count: 1

Selected Inconsistency: *eat(cake)* mutex with *have(cake)* no-op

Neighborhood: Remove *eat(cake)* or *-have(cake)* no-op

Walkplan Example




Inconsistency count: 1

Selected Inconsistency: unsupported fact *have(cake)*

Neighborhood: add *bake(cake)* action or no-op

Walkplan Example




Inconsistency count: 2

Selected Inconsistency: unsupported fact $\neg have(cake)$

Neighborhood: add $bake(cake)$ action

Walkplan Example




Inconsistency count: 3

Selected Inconsistency: mutex between *eat(cake)* and $\neg eaten(cake)$ no-op

Neighborhood: remove $eat(cake)$ or $\neg eaten(cake)$ no-op

Walkplan Example



Inconsistency count: 2

Selected Inconsistency: mutex between *eat(cake)* and *have(cake)* no-op

Neighborhood: remove *eat(cake)* or *have(cake)* no-op


```
graph TD
    start([start]) --> read[read]
    read --> write[write]
```



Selected Inconsistency: unsupported fact *have(cake)*

Neighborhood: remove *eat(cake)* or add *have(cake)* no-op

Walkplan Example




Inconsistency count: 1

Selected Inconsistency: unsupported fact *eaten(cake)*

Neighborhood: add *eat(cake)* or add *eaten(cake)* no-op

Walkplan Example



Inconsistency count: 0
Return finished plan

Takeaways

- ▶ Local search through the space of planning graphs can be more effective and scalable than depth-first search.
- ▶ Calculating persisting mutexes is more efficient.
- ▶ Local search allows for plan refinement if you have an initial plan to start with.