

Anytime Search in Dynamic Graphs

Paper by:
M. Likhachev
D. Ferguson
G. Gordon
A. Stentz
S. Thrun

Presentation by:
Max Pflueger

January 19, 2011

Previous search algorithms

- A^* (optimal)
- weighted A^* (approximate within factor of optimal)
- Anytime A^*
 - find an approximate solution, then improve it by continuing to run optimal A^*
- real-time or agent-centered searches
 - initially restrict state space and expand as time permits
 - provide no guarantees
- LPA* and D*
 - handles dynamic graphs, but not anytime

Generalization of A*

- introduces v -values
 - define a state as *inconsistent* if $v(s) \neq g(s)$
 - specifically, if $v(s) > g(s)$ *overconsistent*, if $v(s) < g(s)$ *underconsistent*

The pseudocode below assumes the following :

(1) v -values of all states are set to ∞ , $g(s_{\text{start}}) = 0$ and the g -values of the rest of the states are set to ∞ (the initialization can also occur whenever `ComputePath` encounters new states);

(2) $OPEN = \{s_{\text{start}}\}$.

```
1 procedure ComputePath()
2 while( $s_{\text{goal}}$  is not expanded)
3   remove  $s$  with the smallest  $f(s)$  from  $OPEN$ ;
4    $v(s) = g(s)$ ;
5   for each successor  $s'$  of  $s$ 
6     if  $g(s') > g(s) + c(s, s')$ 
7        $g(s') = g(s) + c(s, s')$ ;
8     insert/update  $s'$  in  $OPEN$  with  $f(s') = g(s') + h(s')$ ;
```

Fig. 4. A* Search: `ComputePath` function with v -values

Generalization of A*

- allow states to be expanded in the order of a key function

for any two states $s, s' \in S$ such that $c^(s', s_{\text{goal}}) < \infty$, $v(s') \geq g(s')$, $v(s) > g(s)$ and $g(s') > g(s) + \epsilon * c^*(s, s')$, it must hold that $\text{key}(s') > \text{key}(s)$*

Fig. 5. key restriction

Generalization of A*

- restrict expansion of states to only once

The pseudocode below assumes the following:


- (1) the function $\text{key}(s)$ satisfies the restriction in Figure 5;
- (2) v - and g -values of all states are initialized in such a way that $v(s) \geq g(s) = \min_{s' \in \text{pred}(s)} (v(s') + c(s', s)) \forall s \neq s_{\text{start}}$ and $v(s_{\text{start}}) \geq g(s_{\text{start}}) = 0$ (the initialization can also occur whenever `ComputePath` encounters new states);
- (3) $CLOSED = \emptyset$ and $OPEN$ contains the overconsistent states (i.e., states s whose $v(s) > g(s)$).

```
1 procedure ComputePath()  
2 while( $\text{key}(s_{\text{goal}}) > \min_{s \in OPEN} (\text{key}(s))$ )  
3   remove  $s$  with the smallest  $\text{key}(s)$  from  $OPEN$ ;  
4    $v(s) = g(s)$ ;  $CLOSED = CLOSED \cup \{s\}$ ;  
5   for each successor  $s'$  of  $s$   
6     if  $g(s') > g(s) + c(s, s')$   
7        $g(s') = g(s) + c(s, s')$ ;  
8       if  $s' \notin CLOSED$   
9         insert/update  $s'$  in  $OPEN$  with  $\text{key}(s')$ ;
```


Fig. 6. A* search with a generalized priority function and generalized overconsistent initialization: `ComputePath` function

Anytime Repairing A* (ARA*)


- Repeat weighted A* with shrinking suboptimality bound
- Reuse results from previous iterations
- No state expanded more than once per iteration



initial search ($\epsilon = 2.5$)



second search ($\epsilon = 1.5$)



third search ($\epsilon = 1.0$)

Fig. 10. ARA* search

ARA*

```
1 procedure ComputePath()
2 while(key( $s_{goal}$ ) >  $\min_{s \in OPEN}(\text{key}(s))$ )
3   remove  $s$  with the smallest key( $s$ ) from OPEN;
4    $v(s) = g(s)$ ;  $CLOSED = CLOSED \cup \{s\}$ ;
5   for each successor  $s'$  of  $s$ 
6     if  $s'$  was never visited by ARA* before then
7        $v(s') = g(s') = \infty$ ;
8       if  $g(s') > g(s) + c(s, s')$ 
9          $g(s') = g(s) + c(s, s')$ ;
10      if  $s' \notin CLOSED$ 
11        insert/update  $s'$  in OPEN with key( $s'$ );
12      else
13        insert  $s'$  into INCONS;
```

Fig. 8. ARA*: ComputePath function. ARA* specific changes as compared with A* search as formulated in Figure 6 are shown in bold.

The pseudocode below assumes the following:


(1) heuristics are consistent: $h(s) \leq c(s, s') + h(s')$ for any successor s' of s if $s \neq s_{goal}$ and $h(s) = 0$ if $s = s_{goal}$.

```
1 procedure key( $s$ )
2 return  $g(s) + \epsilon * h(s)$ ;


3 procedure Main()
4  $g(s_{goal}) = v(s_{goal}) = \infty$ ;  $v(s_{start}) = \infty$ ;
5  $g(s_{start}) = 0$ ;  $OPEN = CLOSED = INCONS = \emptyset$ ;
6 insert  $s_{start}$  into OPEN with key( $s_{start}$ );
7 ComputePath();
8 publish current  $\epsilon$ -suboptimal solution;
9 while  $\epsilon > 1$ 
10   decrease  $\epsilon$ ;
11   Move states from INCONS into OPEN;
12   Update the priorities for all  $s \in OPEN$  according to key( $s$ );
13    $CLOSED = \emptyset$ ;
14   ComputePath();
15   publish current  $\epsilon$ -suboptimal solution;
```

Fig. 9. ARA*: key and Main functions


ARA* Experimental comparison



(a) 6D arm trajectory for $\epsilon = 3$



(b) uniform cost results



(c) non-uniform cost results

Fig. 11. 6D robot arm experiments. Anytime A* is the algorithm presented in [16].

Anytime Dynamic A* (AD*)

- apply the earlier technique for anytime planning to LPA*
- LPA* can deal with changes to state transition costs
 - earlier framework allows it to handle anytime planning simultaneously

Examples

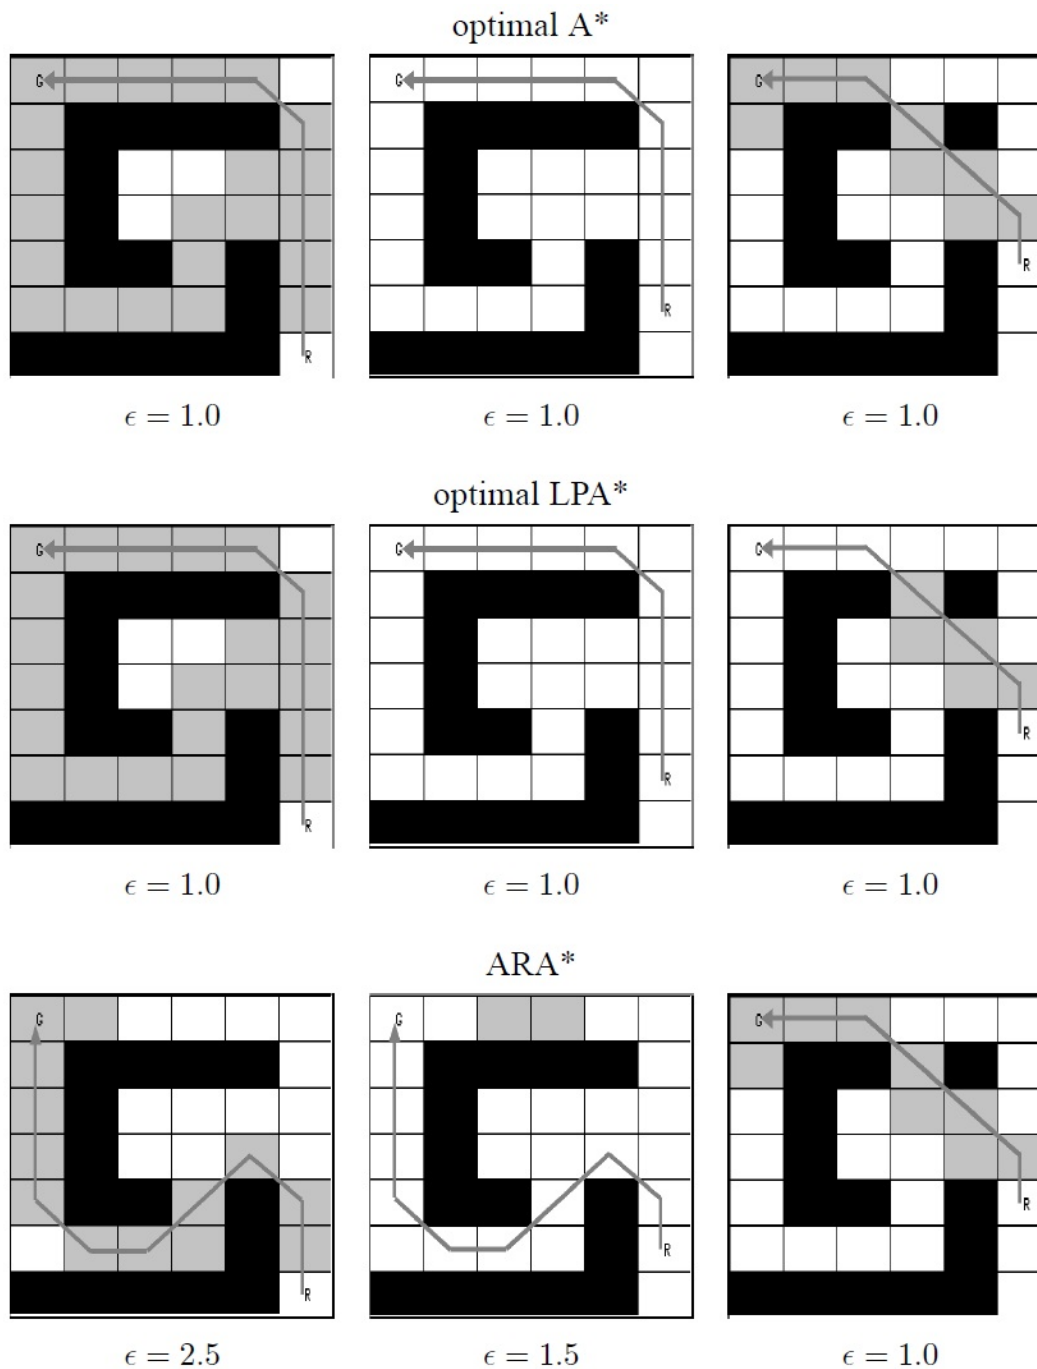


Fig. 20. An example of planning with optimal A*, optimal LPA*, and ARA*. Each algorithm directed its searches from the goal state G to the agent state R. The states expanded by the algorithms are shown in grey. Note that after the third planning episode each of the algorithms can guarantee solution optimality ($\epsilon = 1.0$).

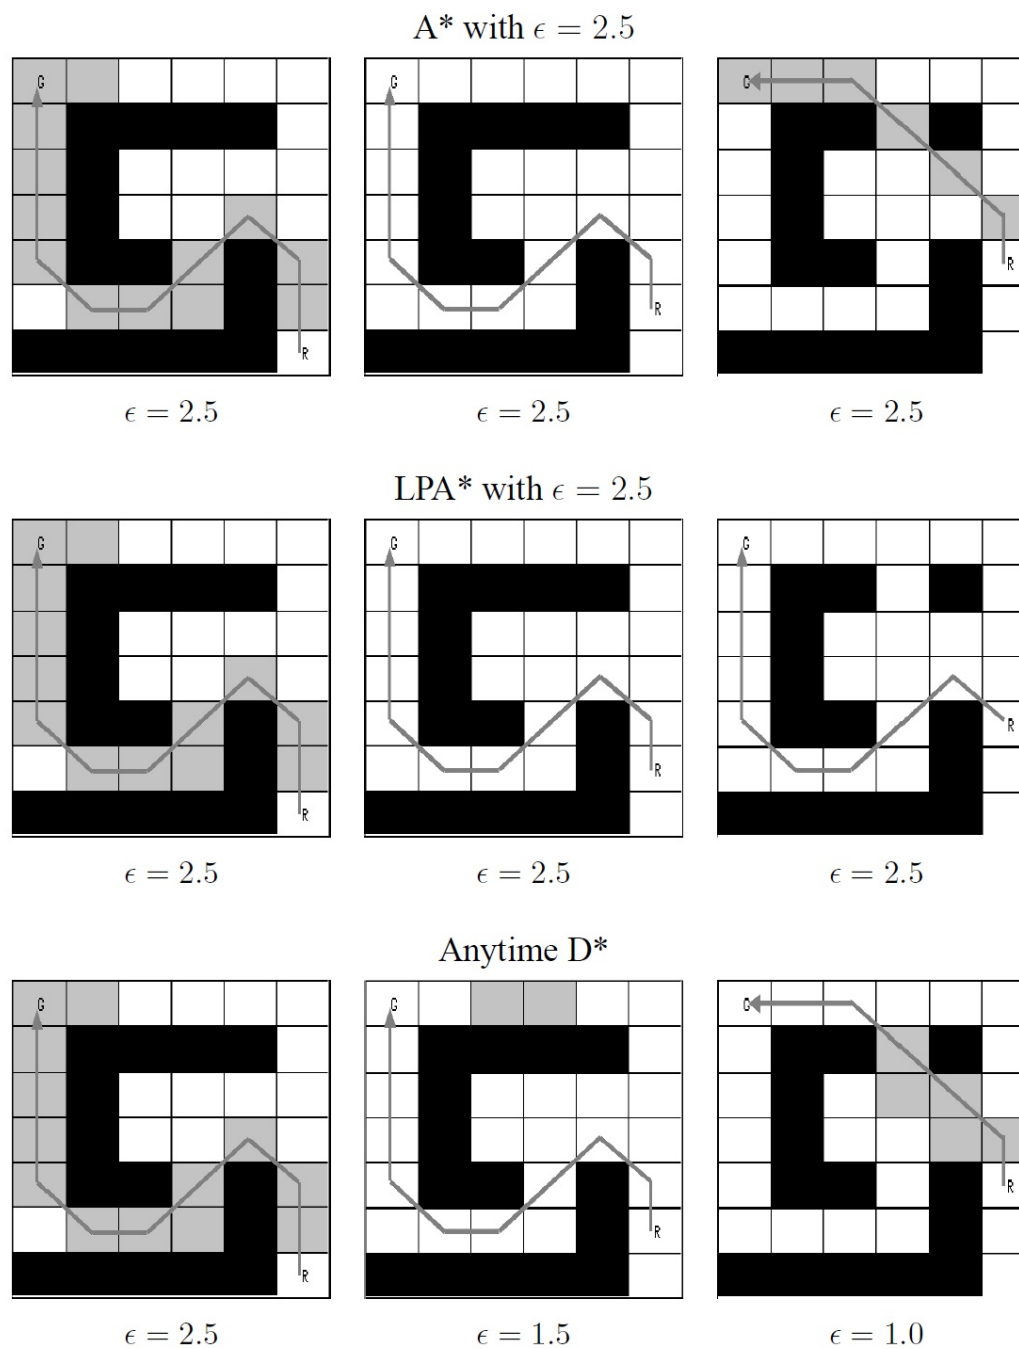


Fig. 21. An example of planning with A* with an inflation factor $\epsilon = 2.5$, LPA* with an inflation factor $\epsilon = 2.5$, and AD*. Each algorithm directed its searches from the goal state G to the agent state R. The states expanded by the algorithms are shown in grey. Note that after the third planning episode only Anytime D* can guarantee solution optimality ($\epsilon = 1.0$).

AD*

```
1 procedure UpdateSetMembership(s)
2 if ( $v(s) \neq g(s)$ )
3   if ( $s \notin CLOSED$ ) insert/update s in OPEN with key(s);
4   else if ( $s \notin INCONS$ ) insert s into INCONS;
5 else
6   if ( $s \in OPEN$ ) remove s from OPEN;
7   else if ( $s \in INCONS$ ) remove s from INCONS;

8 procedure ComputePath()
9 while( $key(s_{goal}) > \min_{s \in OPEN}(key(s))$  OR  $v(s_{goal}) < g(s_{goal})$ )
10  remove s with the smallest key(s) from OPEN;
11  if  $v(s) > g(s)$ 
12     $v(s) = g(s)$ ;  $CLOSED = CLOSED \cup \{s\}$ ;
13    for each successor s' of s
14      if s' was never visited by AD* before then
15         $v(s') = g(s') = \infty$ ;  $bp(s') = \text{null}$ ;
16        if  $g(s') > g(s) + c(s, s')$ 
17           $bp(s') = s$ ;
18           $g(s') = g(bp(s')) + c(bp(s'), s')$ ; UpdateSetMembership(s');
19      else //propagating underconsistency
20         $v(s) = \infty$ ; UpdateSetMembership(s);
21        for each successor s' of s
22          if s' was never visited by AD* before then
23             $v(s') = g(s') = \infty$ ;  $bp(s') = \text{null}$ ;
24            if  $bp(s') = s$ 
25               $bp(s') = \arg \min_{s'' \in pred(s')} v(s'') + c(s'', s')$ ;
26               $g(s') = v(bp(s')) + c(bp(s'), s')$ ; UpdateSetMembership(s');
```

Fig. 18. Anytime D*: ComputePath function. The changes as compared with the ComputePath described in section 5.2 are shown in bold.

AD*

The pseudocode below assumes the following:


(1) heuristics are consistent: $h(s) \leq c(s, s') + h(s')$ for any successor s' of s if $s \neq s_{\text{goal}}$ and $h(s) = 0$ if $s = s_{\text{goal}}$.

```
1 procedure key( $s$ )
2   if ( $v(s) \geq g(s)$ )
3     return [ $g(s) + \epsilon * h(s); g(s)$ ];
4   else
5     return [ $v(s) + h(s); v(s)$ ];


6 procedure Main()
7    $g(s_{\text{goal}}) = v(s_{\text{goal}}) = \infty; v(s_{\text{start}}) = \infty; bp(s_{\text{goal}}) = bp(s_{\text{start}}) = \text{null};$ 
8    $g(s_{\text{start}}) = 0; OPEN = CLOSED = INCONS = \emptyset; \epsilon = \epsilon_0;$ 
9   insert  $s_{\text{start}}$  into  $OPEN$  with key( $s_{\text{start}}$ );
10  forever
11    ComputePath();
12    publish  $\epsilon$ -suboptimal solution;
13    if  $\epsilon = 1$ 
14      wait for changes in edge costs;
15      for all directed edges  $(u, v)$  with changed edge costs
16        update the edge cost  $c(u, v)$ ;
17        if ( $v \neq s_{\text{start}}$  AND  $v$  was visited by AD* before)
18           $bp(v) = \arg \min_{s'' \in \text{pred}(v)} v(s'') + c(s'', v);$ 
19           $g(v) = v(bp(v)) + c(bp(v), v);$  UpdateSetMembership( $v$ );
20    if significant edge cost changes were observed
21      increase  $\epsilon$  or re-plan from scratch (i.e., re-execute Main function);
22    else if  $\epsilon > 1$ 
23      decrease  $\epsilon$ ;
24    Move states from  $INCONS$  into  $OPEN$ ;
25    Update the priorities for all  $s \in OPEN$  according to key( $s$ );
26     $CLOSED = \emptyset;$ 
```

Fig. 19. Anytime D*: key and Main functions

AD* experimental comparison



optimal solution



end-effector trajectory






Fig. 23. Environment used in our second experiment, along with the optimal solution and the end-effector trajectory (without any dynamic obstacles). Also shown are the solution cost of the path traversed and the number of states expanded by each of the three algorithms compared.

Application to Mobile Robots



(a) optimal 4D search with A*
after 25 secs

(b) 4D search with AD*
after 0.6 secs ($\epsilon = 2.5$)

(c) 4D search with AD*
after 25 secs ($\epsilon = 1.0$)

Fig. 27. The comparison of planning with A* and Anytime D* for outdoor robot navigation (cross shows the position of the robot). Both 4D searches used 2D search to compute heuristics when necessary as described in the text. In this example no information inconsistent with the initial map was observed during execution (i.e. no edge cost changes occurred).