

Koenig and Likhachev, "D*-Lite"

- Motivation
 - Path planning is a problem -- especially if we want it to be optimal
- Lifelong Planning A* (LPA*)
 - A*
 - g-value: path cost function
 - admissible heuristics: never overestimates
 - in this paper $h(s) = \max\{|x_s - x_g|, |y_s - y_g|\}$ (maximum of the x distance and y distance)
 - $f(s) = g(s) + h(s)$ --> how the algorithm decides the order in which to visit nodes
 - A* finds minimal cost from given start to goal --> but what if something changes along the way?
 - A* has to throw everything away and re-plan
 - these approaches are iterative --> reuse previous results to compute a new path in the changed environment, while preserving the guarantee of finding shortest path
 - LPA* --> incremental version of A*
 - based on the assumption that there are typically only small changes
 - maintains two estimates of the start distance
 - g-value
 - distance from start to current position
 - rhs-value
 - estimate of minimal distance to node (lookahead value), should be equal to g-value
 - updates the g-values of the pervious search
 - $rhs(s) = 0$ if $s = s_start$ and $\min_{s' \text{ predecessor of } s} (g(s') + c(s', s))$ otherwise
 - a vertex is consistent if its g-value is its rhs value
 - Algorithm:
 - initialize g-values with infinity
 - expand cell with smallest f-value
 - if $g(s) > rhs(s)$ -> replace. Otherwise, set to infinity
 - === CHANGE IN GRAPH ===
 - recalculates the rhs values of potentially-affected cells and checks for consistency
 - remove them or add to the priority queue
 - repeat until shortest path is found
 - Question: when do you calculate rhs
 - in LPA* start and end values are constant
 - in D* lite, the start value is changing
 - not everything has to be consistent, only the stuff that's at the top of the priority queue -- this may not be just along the optimal path --

when there's a change the rhs value changes for the adjacent nodes, which will make them inconsistent, and then the change will propagate as far as necessary for the optimal path

- this may be the same as the v values from the other paper
 - Properties of LPA*
 - similar to A*
 - terminates correctly
 - efficient
 - Example: <http://idm-lab.org/project-a.html> (fun to play with!)
- D* Lite
 - solves goal-directed navigation problems in unknown terrain
 - moving starting point
 - searches from goal to start (backwards)
 - first step, same as LPA*, then move start position along the path, checking for changes in environment
 - because position is moving, you would have to reorder priority queue
 - to avoid this, $h(s, s')$ must be nonnegative, and satisfy $h(s, s') \leq c^*(s, s')$ and $h(s, s'') \leq h(s, s') + h(s', s'')$
 - this is because we have nodes that we expanded, but we don't need to expand anymore, so heuristic value goes from frontier to robot location, so in queue heuristic value will get smaller when the robot moves. So, although the values change, the ordering on the queue doesn't change.
 - seems like there should be a subtraction here -- where?
 - doesn't happen in the algorithm...
 - it would be a lower bound if you subtracted, but we don't need to, because the ordering is maintained
 - but need to keep track of this: keep the K_m value and add it back in later
 - this is not the same as D*
 - digression -- what is D*
 - does the same dynamic robot moving thing and reusing previous work
 - this is a suggested simpler version which is easier to analyze
 - this is used more often because it's easier to implement and analyze and works better
 - does the same thing but better
 - all the optimizations (like heap reordering) don't appear in D*
 - in paper it says that heap reordering uses a method derived from D*, but there's no explanation...
 - guarantee termination and that shortest path will be found -- intuition - if g-values are consistent then shortest path will be found, and will terminate since only expands each node at most once.
 - digression -- properties of algorithms
 - termination is important too
 - this algorithm has the same guarantees of optimality and completeness as older algorithms
 - Experimental Results
 - compare D* to D*lite experimentally

- D* lite always performs at least as well as D*, but often better
- this is because of the smart reordering
- also show that D* lite expands fewer cells than regular A*
- Questions/Comments
 - note: A* has a nice property of expanding the optimal number of nodes - no nodes that were expanded that should not have been. D* and D*lite don't have this, so there could be a D*lite-ER which could expand even fewer nodes and be optimal
 - D* is a greedy algorithm
 - in an unknown environment D* might not do well by moving greedily towards the goal
 - D* doesn't reason about the probability of obstacles, so it can't take "safer" paths when there might be an obstacle near the goal
 - stochastic methods do this better
 - belief A* changes the heuristic
 - but not optimal if you have information about your environment, to get optimal you'd have to use a POMDP for example
 - there are constraints that aren't easy to fit into this
 - but does this ruin it for dynamic environments?
 - no, because optimality is not as important as termination in many real-world scenarios
 - presence of obstacles -> can get info from old maps, or multi-robot systems, or the environment is changing fast
 - have also looked at comparisons of D* and the perfect information case
 - naively a bound of $|P(D^*)| \leq v^2 |P(\text{opt})|$ (can tighten to $v^{3/2}$) $v = \#$ of nodes
 - show cases that reveal the problems in the greedy approach
 - Guarav recommends implementing algorithms in order to understand what's going on
 - lots of bookkeeping work

Likhachev et al., "Anytime Search in Dynamic Graphs"

- This was implemented on a version of the DARPA grand challenge -- getting close to the state-of-the-art
- Previous search algorithms
 - A*
 - optimal
 - weighted A*
 - guarantees on suboptimality
 - anytime A*
 - find approximate solution, improve it
 - real-time/agent-centered searches
 - initially restrict state space and expand as time permits
 - no guarantees
 - LPA* and D*
 - handle dynamic graphs, but not anytime
- Want guarantees on suboptimality, dynamic, anytime algorithm
- Generalization of A*
 - introduce v-values - similar to rhs values - if the heuristic is admissible, $v(s) = g(s)$.
 - state is inconsistent if $v(s) \neq g(s)$, if $v(s) > g(s)$ overconsistent, if $v(s) < g(s)$, underconsistent

- allow states to be sorted by a key function
 - properties: if a state can reach the goal, if the state, s' and s are overconsistent and $g(s') > g(s) + \epsilon \times c^*(s, s')$, then $key(s') > key(s)$
- allow A^* to start with values already assigned to states, and force to expand states at most once (by sticking them in a set “CLOSED”)
- Anytime Repairing A^* (ARA*)
 - repeat weighted A^* with shrinking suboptimality bound
 - reuse results from previous iterations
 - no state expanded more than once per iteration
 - track inconsistent states which have been added to closed set, because the algorithm will have to check those out again.
 - if you run out of time with highest solution -- no path! so make sure first epsilon is chosen to be big enough to get some solution quickly
 - still no dynamic planning
 - experimental comparison - anytime A^* , A^* , ARA*
- Anytime Dynamic A^* (AD*)
 - apply the earlier technique for anytime planning to LPA*
 - intuitively, we want to only search the new states when the graph changes, while still being able to tighten epsilon constraints.
 - use backpointers to track which nodes depend on nodes whose values have changed
 - propagate underconsistency
 - if we hit an underconsistent state, set its v value to infinity, then follow its
 - if a state is inconsistent, and its successor hasn't been visited before, set v and g to infinity, backpointer to null
 - given changed computePath() function, main function is basically the same as LPA*, but checking for changes and looping forever
 - NB: the “if significant edge cost changes observed” line is the key to this whole algorithm, because the problem is significance rather than number or magnitude. A small change might change everything -- want to throw everything away sometimes, but hard to decide when to do so. (line 20 of algorithm) Also tough to decide when to change epsilon versus when to leave it the same.
 - compare performance as probability of encountering a random obstacle increases
 - discussion: what would happen if increased beyond 20? why didn't they show this?
 - applications to mobile robot navigation
 - do 4d navigation to control orientation and velocity because of robot limitation
 - use 2d A^* as the heuristic for the 4d search
 - robot using AD* starts moving after 0.6 seconds, with A^* moves after 11s
 - AD* robot still continues refining plan as it moves
 - also guaranteed not to be worse than 2.5 times the optimum!
 - see fig. 29 for an example with no initial information
 - used on the DARPA urban challenge (3rd challenge)
 - <http://www.youtube.com/watch?v=SQFEmR50HAK>
 - digression: first darpa challenge -> no one won
 - motorcycle in urban challenge fell over - haha
 - <http://www.youtube.com/watch?v=6CYGT97i8qU>

Additional Info on DARPA paper

- Lattice planner
 - non polynomial constraints for cars
 - limit possible states into which car can transition (constant branching factor)
 - can change the resolution of the lattice
- Apply anytime incremental search to the lattice
- Application is parking the vehicle
- 2 heuristics, mechanism-constrained and environmentally constrained -- take the maximum of the two
- taking maximum of the two does order of magnitude better than each individually
- does pretty well - real-world application of planning
- ignores dynamics of the car
 - some later work does
 - example: what speeds you can go at depending on turn angle
 - this does x, y, v where v is max forward or max backwards velocity
- really would do continuous control on the low-level stuff and do planning on high-level goals
- cool video: stanford robotic car slides into parking space:
 - <http://singularityhub.com/2010/05/12/stanfords-robot-car-slides-into-parking-spot-like-a-badass-video/>